TRIUMF Site Report
Tier-1

Denice Deatrich

HEPIX

October 24, 2011 Vancouver

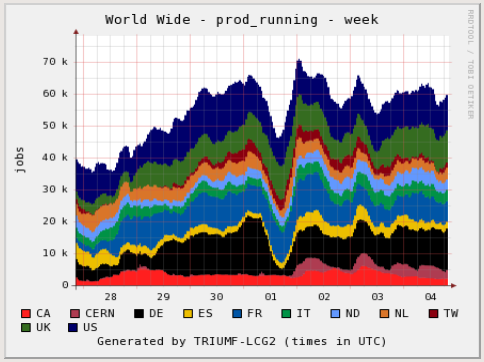

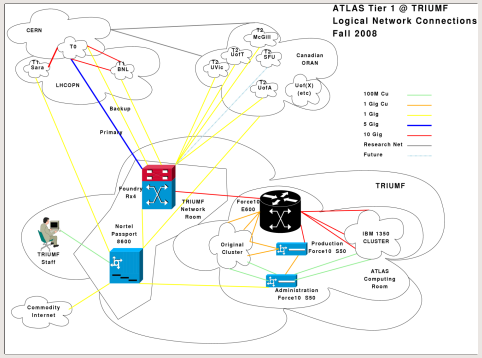

Outline

- TRIUMF Tier-1
  - Review personnel and operational status
  - Review current cluster & infrastructure configuration
  - Look at some topical issues
  - Peek at the procurement plans for this year
- TRIUMF Site-Wide
  - Highlight some recent TRIUMF networking changes

Operations I

- Systems support & operations personnel:
  - Group leader, 5 experts (systems, grid, databases, network, storage), 1 hardware specialist
  - 24x7 on-call rotation shared by Systems experts and group leader
- User support personnel:
  - 3 experts - ATLAS software, Grid applications, Tier-3 tools and user analysis in the Canadian cloud
- Data centre current capacity and server numbers:
  - 1210 cores (232 worker nodes, all blades)
  - 2.1 PB disk (usable), 5.5 PB of tape (recently upgraded from 1.6 PB)
  - 97 servers (storage, middle ware, admin, standby and test servers). 30 of these are virtual. Note that the resources used by the user support personnel are not counted here.

Operations II

- Redundant configurations for effective high availability
  - UPS, redundant cooling units, hardware redundancy in dual power supplies, redundant networking, raided disks, ...
  - Critical service redundancy using virtualization with configured virtual standby servers (e.g. site-bdii / site-bdii-sb)
  - Nagios sensors for email & SMS notification of T1 personnel.
  - Automated shutdown of some systems for critical temperature events (especially worker nodes)
  - System log monitoring and cron job alerts for hardware issues not yet covered by Nagios
  - Site LAN/WAN is monitored 24x7 (expert paging since 2007)
  - Grid services monitored by WLCG (SAM, GGUS alarms)
  - Redundancy in personnel (holidays, illness, etc)

Tier-1 Availability

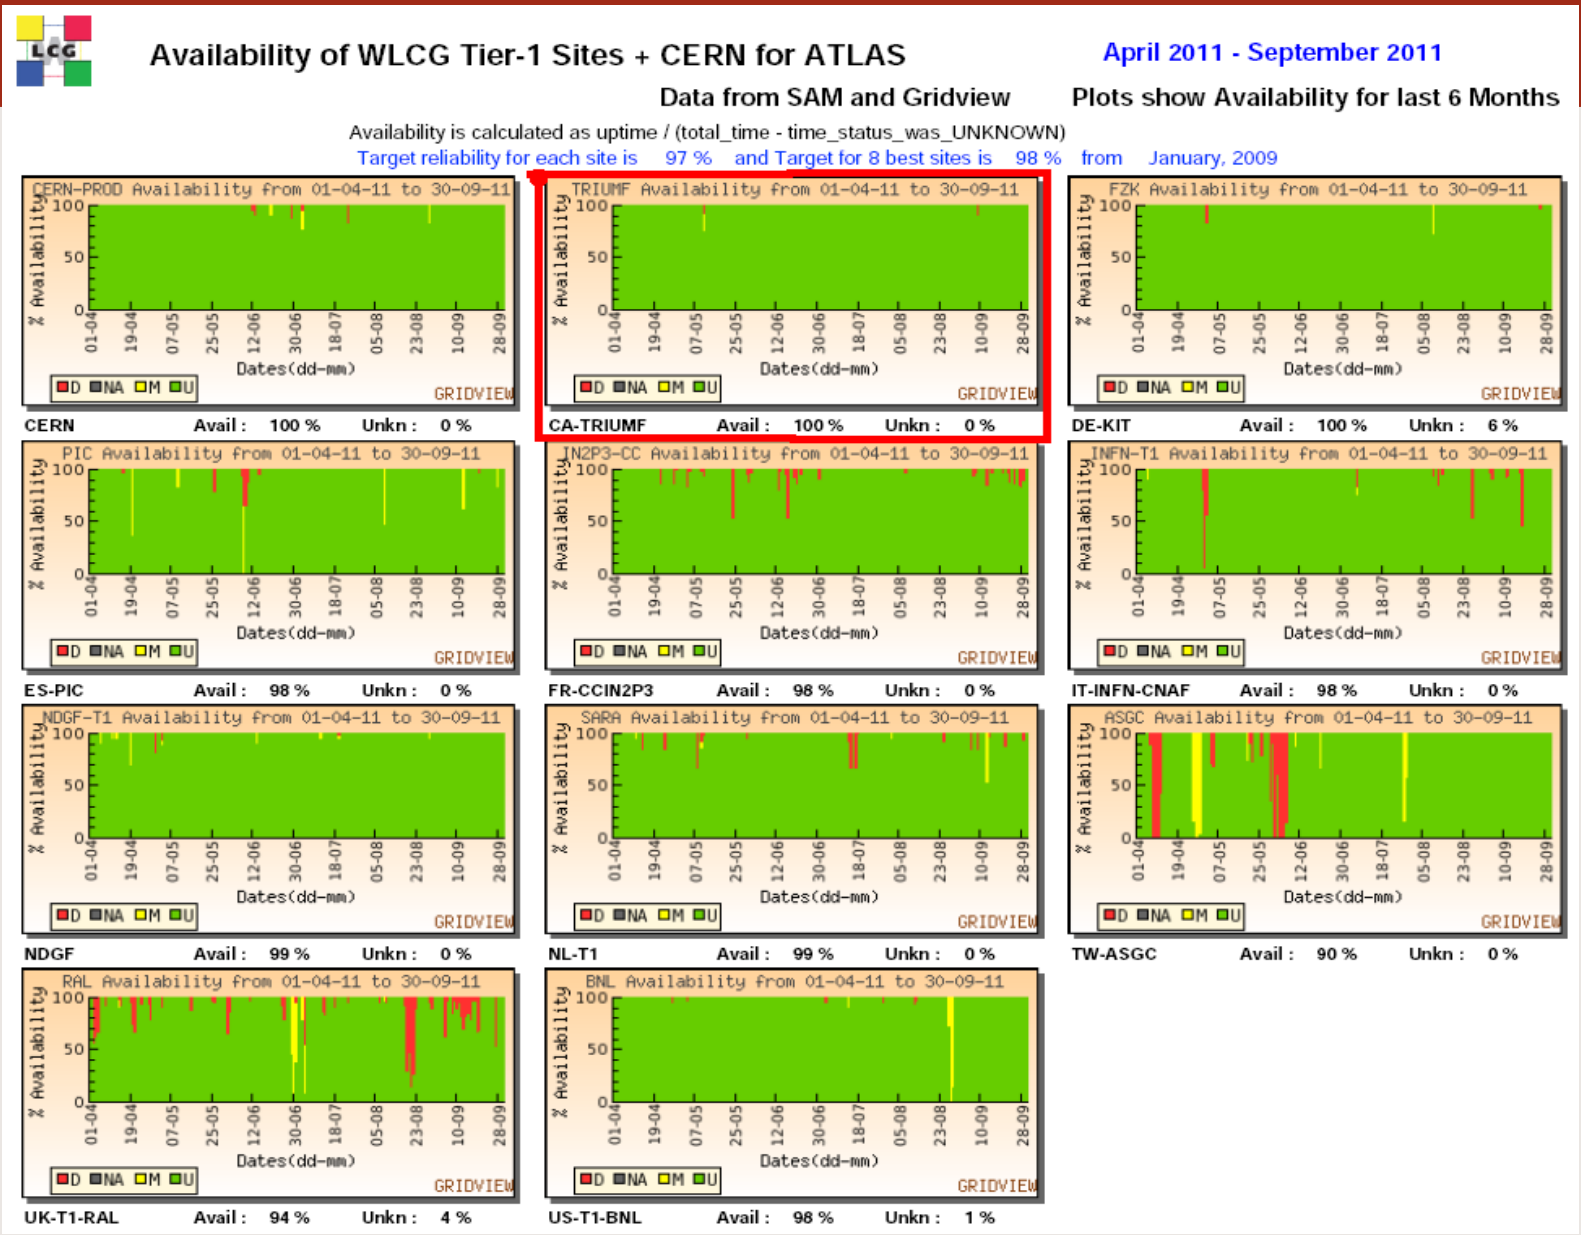

Cluster overview

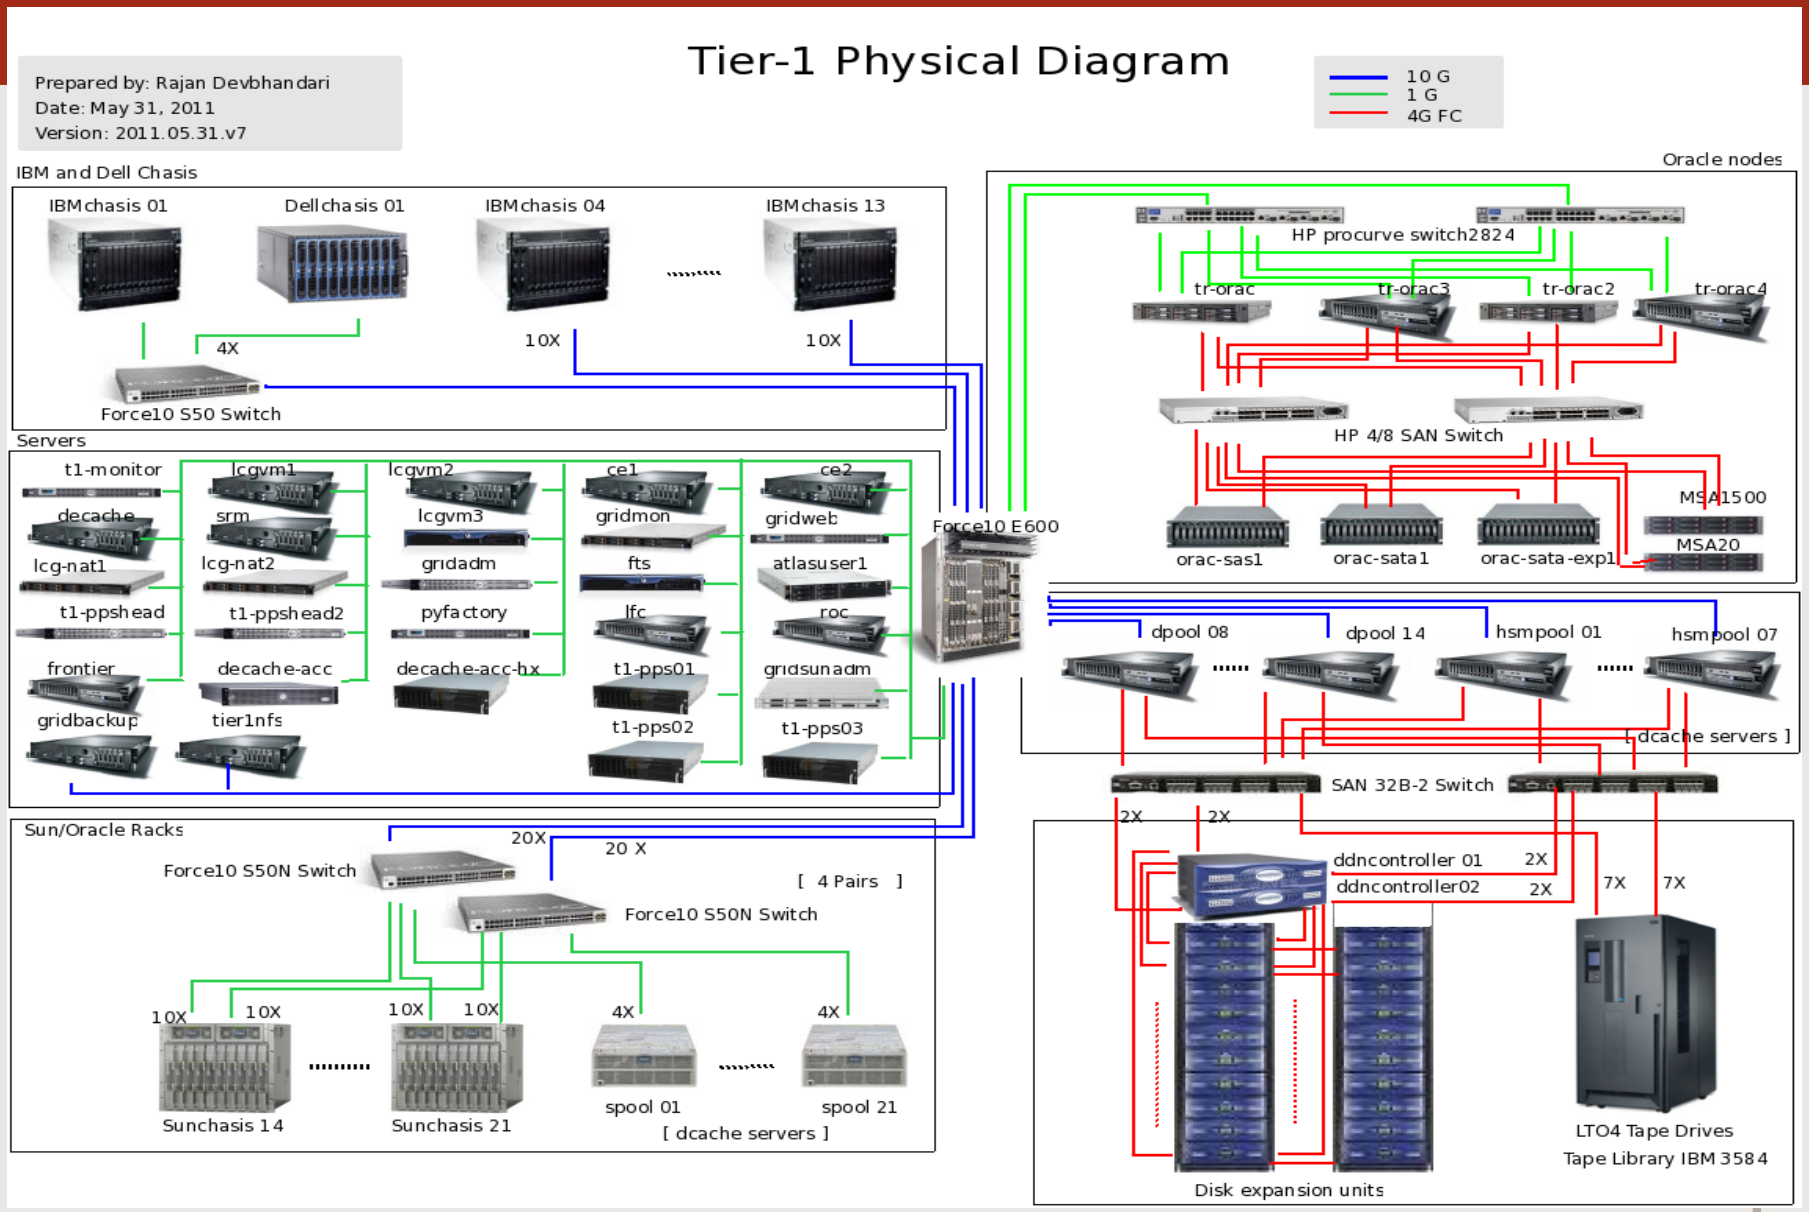

Infrastructure

- **Limited floor space**: 43' x 22'
- **No false floor**
- **Rack optimization:**
  - high density solution
  - hot & cold configuration
- **Power estimate:**
  ~0.4 MW (by 2012)
  (including cooling)

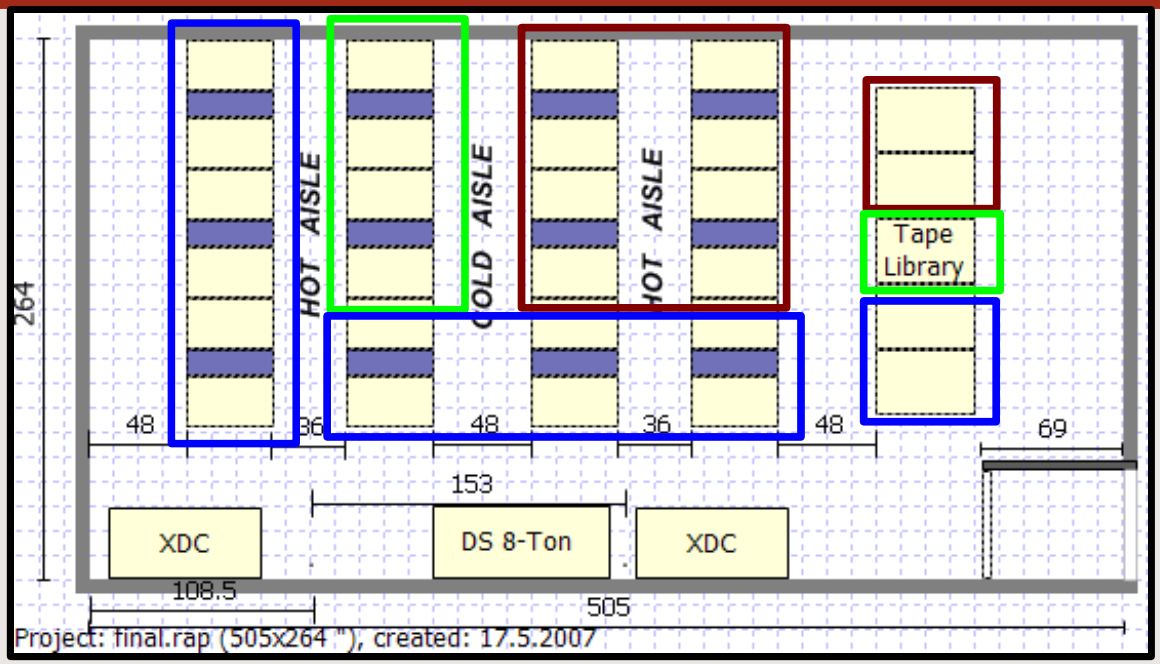

- **Cooling solution:** Liebert XD system (very flexible)
  340 kW total capacity (~35% used)
- **UPS:** 225 kVA (in the future CPU racks on regular power or expand UPS capacity) (~55% used).
  (no diesel backup except for core network)

2007+2008 (2013/2014 ?)

2009+2010 (higher density / LHC delay)

2011+

(slide courtesy of Réda Tafirout)

- **Need to explore future expansion scenarios for 2015 and beyond (Tier-1 will be a 10 years old facility...).**

Operating system status

- Operating systems and architectures
  - Linux
    - Mostly Linux, all are at SL 5.5-5.7 64-bit (no 32-bit SL5 nodes)
    - 4 Red Hat EL 5.7 (Oracle RACs)
    - 4 SL4 nodes left – 2 management nodes, 1 Atlas software NFS server, and 1 virtual (32-bit) node – should be updated to SL5 by the year's end
  - Solaris
    - 22 nodes – 21 Thors (x4540) and 1 management node - Solaris 10 9/10 (Update 9)
    - very conservative – only updating once or twice a year

Virtualization Status

- 4 Xen servers (will purchase at least one more)
  - 2 servers for general Tier-1 needs
    - dual cpu, dual and quad core, 3 hardware raid-1 disks on each, 10k or 15k SAS
  - 2 servers dedicated to Canadian ROC (prod and standby)
  - guest disks are LVM partitions spread across the available raided disks
  - guest swap disks are sparse (usually 1Gb) files
  - guests either provide production services, or operate as standby for other production servers. We also have a few pre-production (test) guests.
  - have implemented fail-over scripts for prod. <=> standby
- Plan to move from SL5 Xen-based to SL6 KVM-based hosts in 2012

Worker nodes - CVMFS

- CVMFS implementation (current version is 2.0.3)
  - since May we have a cvmfs implementation (ce3.triumf.ca) with 2 blade chassis (28 workers)
  - /home re-partitioned on workers – borrowed 25 GB from /home to create a dedicated partition (/t1cvmfs) for the cvmfs cache.
    CVMFS_QUOTA_LIMIT=22500
  - we reduce the root reserve to 1% on /t1cvmfs (also /home)
  - created an RPM (**t1-cvmfs-config**) to manage the local setup
- Local CVMFS Squid server
  - squid server implemented and configured on our Atlas software NFS server
- Issues
  - So far the main issue is cvmfs updates – the current need to drain the batch queue is a *distinct disadvantage*

Worker node – disk failures

- Disk failure monitoring
  - high disk failure rate in our IBM blades ~ 30/year in past 2 years
  - the problem manifests as partition *read-only* remounts
  - these disks are at least 10 degrees warmer than other Vendor blades. We analyzed disk temperatures in the data centre, looking at the Average Failure Rates for all disks.
  - we implemented a cron job to monitor the partitions every 15 min. It stops or kills pbs_mom to prevent black holes. (Essentially the script tries to touch and remove a file in each partition)

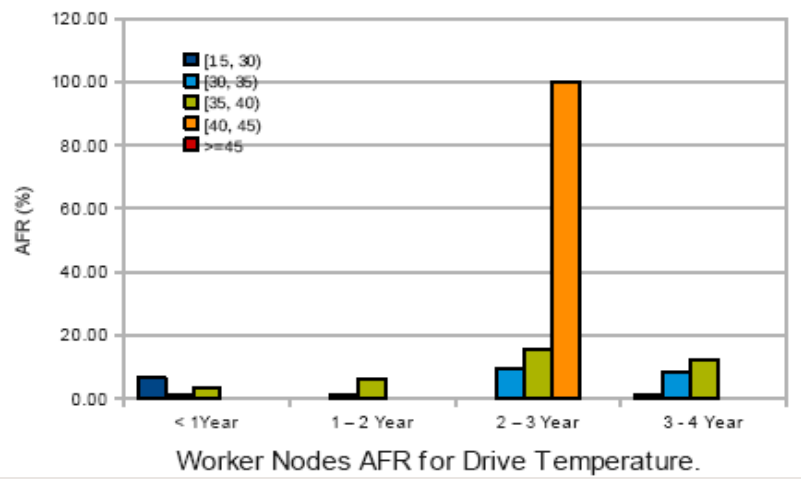

231	231				
Age Group	Total Disk #	< 1Year	1 – 2 Year	2 – 3 Year	3 - 4 Year
[15, 30)	78	6.41	0.00	0.00	1.28
[30, 35)	85	1.18	1.18	9.41	8.24
[35, 40)	65	3.08	6.15	15.38	12.31
[40, 45)	3	0.00	0.00	100.00	0.00
>=45		0.000	0.000	0.000	0.000

(disk failure analysis by Hongyun Sun)

Worker node – lvm stripe-0

- In 2010 we started doing large LVM stripe-0 configurations on worker nodes for /home when 2 disks are available:
  - /dev/sda – swap, /boot, /, /usr, /var, /tmp, /opt (~ 12 GB total)
  - lvm volume group - combined /dev/sda9 + /dev/sdb
  - create 2 lvm stripe-0 partitions for /home and cvmfs. In kickstart use a few tricks to issue the command:
    # lvcreate -i2 ... (will do in kickstart post-install in the future)

    # lvs –segments

```
  LV    VG   Attr   #Str Type    Ssize
  cvmfs vg0  -wi-ao    2 striped  25.00G
  home  vg0  -wi-ao    2 striped 230.00G
```

- Advantage: increases amount of available disk space, augments performance, provides graceful exit on disk failure
- We favour 2-disk worker node configs. in next procurement

2011 Procurement

- Last large purchase to spend balance of our current CFI [1] capital grant
  - RFP goes out immanently ; vendor response deadline will be 3-4 weeks later
    - Disk storage: additional 4 PB usable
    - CPU: additional 3500 cores
  - 2011 purchases (last few months)
    - additional 3.9 PB of tape (LTO5) for our IBM TS3500
      - additional 2 high density frames with 10 LTO 5 drives and 2600 tapes
      - see Simon Liu's presentation on Thursday
    - extra 10g line cards for our Force10 E600 core switch
    - Liebert in-row coolers for new racks
  - will also purchase additional servers as needed in the near future

[1] Canadian Foundation for Innovation

TRIUMF Network Upgrade I

- 2003-2011
  - TRIUMF used CWDM technology to multiplex up to 11 external n/w circuits (mostly ATLAS light-paths) onto two pairs of fiber between TRIUMF and its network providers - BCNet and CANARIE
  - With the exception of the 10Gbps LHCOPN circuit to CERN T0, all other circuits were limited to 1Gbps.
- In April 2011 moved away from CWDM and leased 9 dedicated fibers with ability to go up to 10Gbps on each
  - 1x 10Gbps circuit to BCNet
    - Shared between R&E networks and Commodity networks
  - 8 dedicated fibers into CANARIE's ROADM network at UBC for the LHC/ATLAS T0, T1 and T2 light-paths

(these slides courtesy of Steve MacDonald)

TRIUMF Network Upgrade II

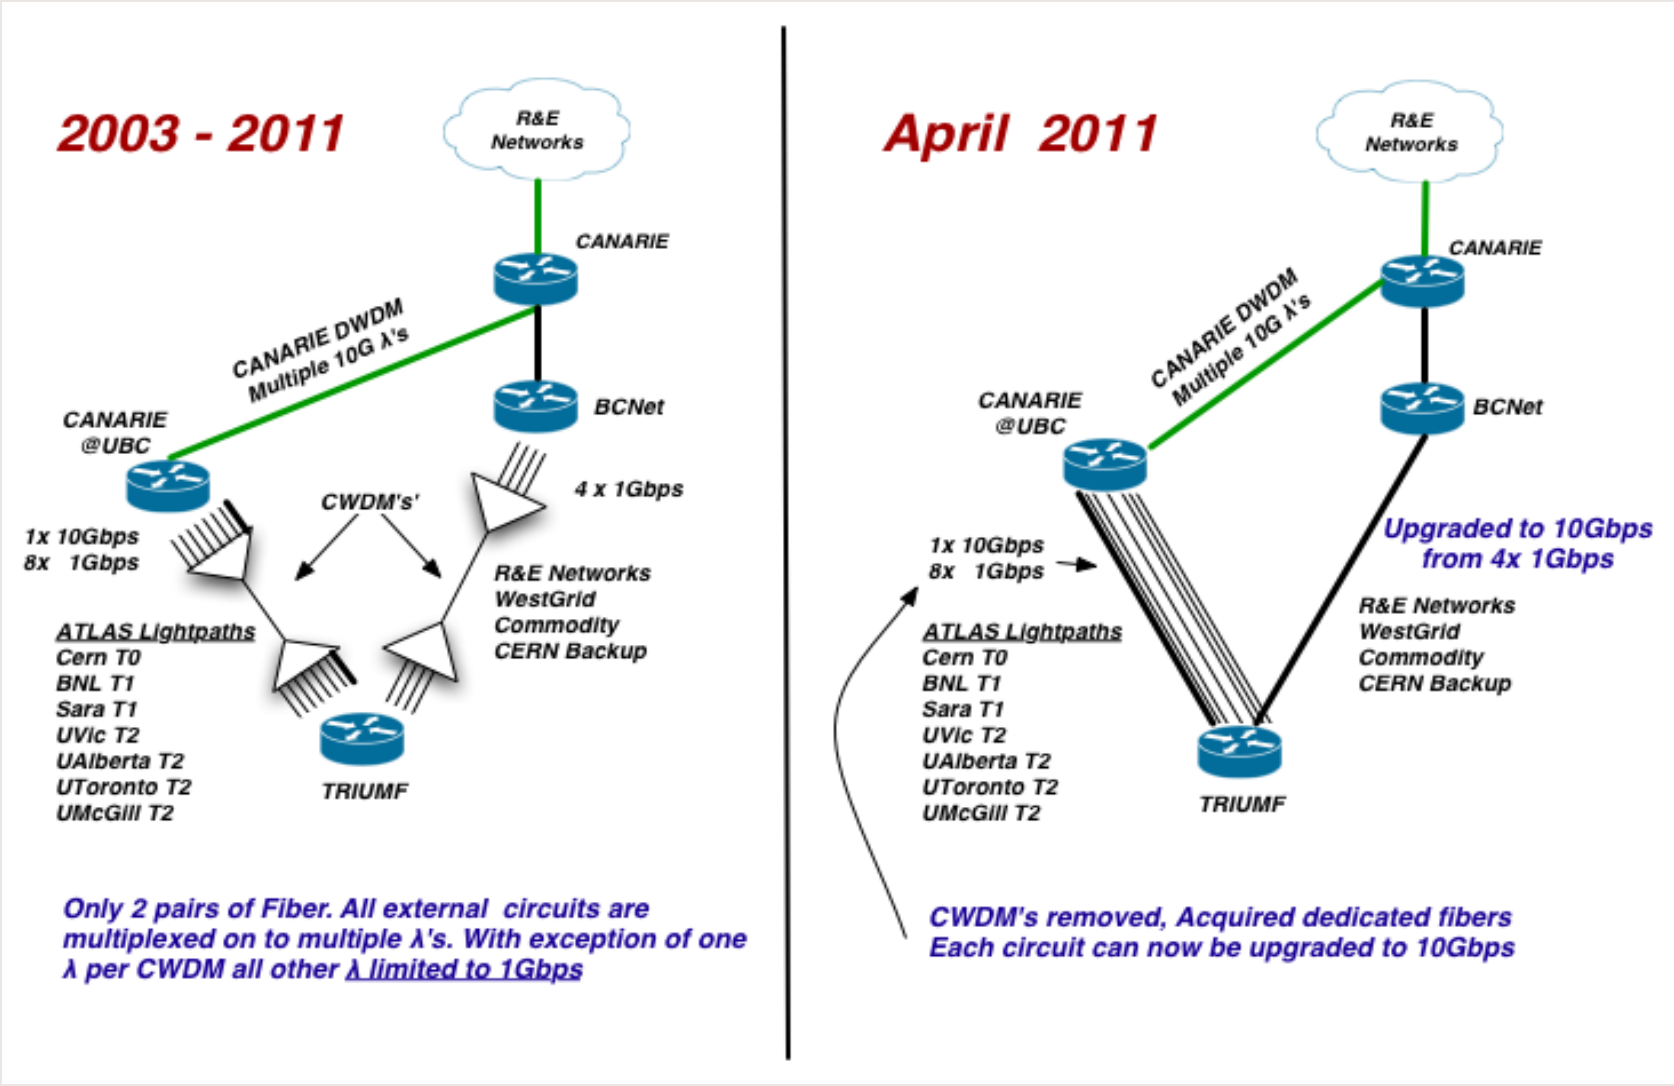

TRIUMF Network Upgrade III

- VoIP Implementation
  - PBX (CS1000M) reached its limit for traditional telephony lines
  - All new buildings now deployed with VoIP
    - Stores/Design office (Sep 2011), next year ARIEL (TRIUMF's new e-linac)
  - Plan to start migrating other areas of TRIUMF that can support VoIP at the edge.
  - All new edge switches PoE capable, currently Avaya/Nortel 5520
  - Telephony becoming responsibility of Networking Group
- Fiber Infrastructure being upgraded
  - Replacing the FDDI grade multi-mode fiber from the 1990's with single-mode fiber to support planned 10Gbps between buildings
- In the process of upgrading network core to 10Gbps

Future Tier-1 Issues

- Data centre infrastructure is sufficient for this TRIUMF planning phase.
- Hardware purchased in 2007 occupies 30% of the room. Its warranty will expire in 2012.
- We are exploring warranty renewal strategies for up to 2014 (e.g. core switch).
- Beyond 2014: a refresh of the data centre will be needed, and will be part of the the next TRIUMF planning phase.

Thank You , Merci !

Acknowledgements:

Canada Foundation for Innovation (CFI)

British Columbia Knowledge Development Funds (BCKDF)

National Research Council (NRC)

National Science & Engineering Council (NSERC)

CANARIE

BCNET

HEPNET

4004 Wesbrook Mall
Vancouver, B.C. Canada V6T 2A3
Tel: 604 222-1047 Fax: 604 222-1074
www.triumf.ca

Cooling I

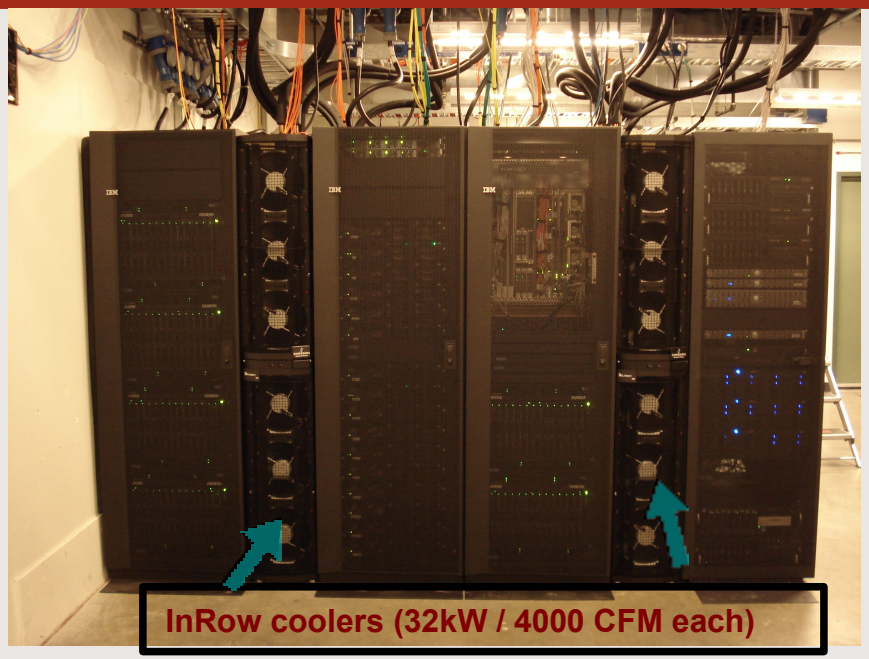

IR picture

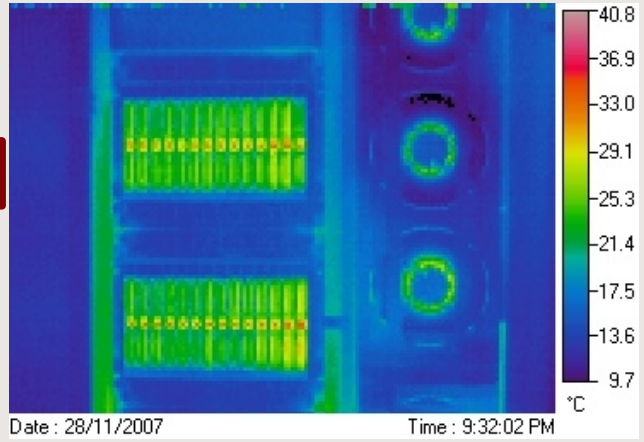

(slide courtesy of Réda Tafirout)

Cooling II

Liebert XDC units: 320 kW cooling capacity total

(heat rejection: refrigerant loop delivered to in-row coolers with air condensers on roof top)

Liebert DS unit: 24 kW cooling (baseline cooling + humidity control)

Redundancy: pumps & compressors

Always above dew point

Total Power consumed:

160 kW (full load)

(slide courtesy of Réda Tafirout)