




CoEPP

ARC Centre of Excellence for
Particle Physics at the Terascale

CoEPP Tier 2 in Melbourne

Sean Crosby
Lucien Boland





Main changes in last year

- Departure of Goncalo Borges
 - head hunted by Uni of Sydney central IT
 - leaving only 2 CoEPP computing staff in 2017
 - Lucien on paternity leave for 4 months
 - Sean running T2 services and CoEPP support alone
- No clear funding for HEP computing in Australia
 - use national facilities wherever possible
 - various small grant applications pending
 - new Australia HEP CoE application to be submitted in 2018
 - in 2018 funding for 1.2 EFT secured to “keep lights on”
- Talks happening to integrate HEP computing into Unimelb HPC
 - Our GPU purchase
 - Petascale campus project
 - Bioinformatics GPU cluster (replacing their BlueGene/Q cluster)
 - Maybe where we also move to....
- LHCONE

T2 Availability/Reliability

Still #1 ATLAS Tier 2 site (a repeat of 2016)

Improved percentages slightly

99.88% Availability (2016 - 99.64%)

99.93% Reliability (2016 - 99.64%)

Site Availability using ATLAS_CRITICAL

From 2016/10 to 2017/10



CPU

no new purchase

acquired hand-me-downs

12 x SGI C2112 (46 servers – 736 cores)

5 x Dell C6145 (10 servers – 640 cores)

DISK

purchased:

8 x Dell R730xd Servers

Intel Xeon E5-2609v3 6C

64 GB RAMM

12 x 8TB 7.2k RPM NLSAS 12Gbps

retired (to our CEPH cluster):

2 x Dell R630 each with 4 x MD1200 PE

In 2017 and 2018, we will provide 1.6% of ATLAS pledge

2018

- 18,318 HS06
- 1.39PB

Our unwritten rule will be broken

- Don't run out of warranty storage
- Next year we'll keep 300TB for an extra year

Compute will be supplemented by using cloud resources more

- Reduction of Tier-3 resources
- Possibly gifted old cloud nodes

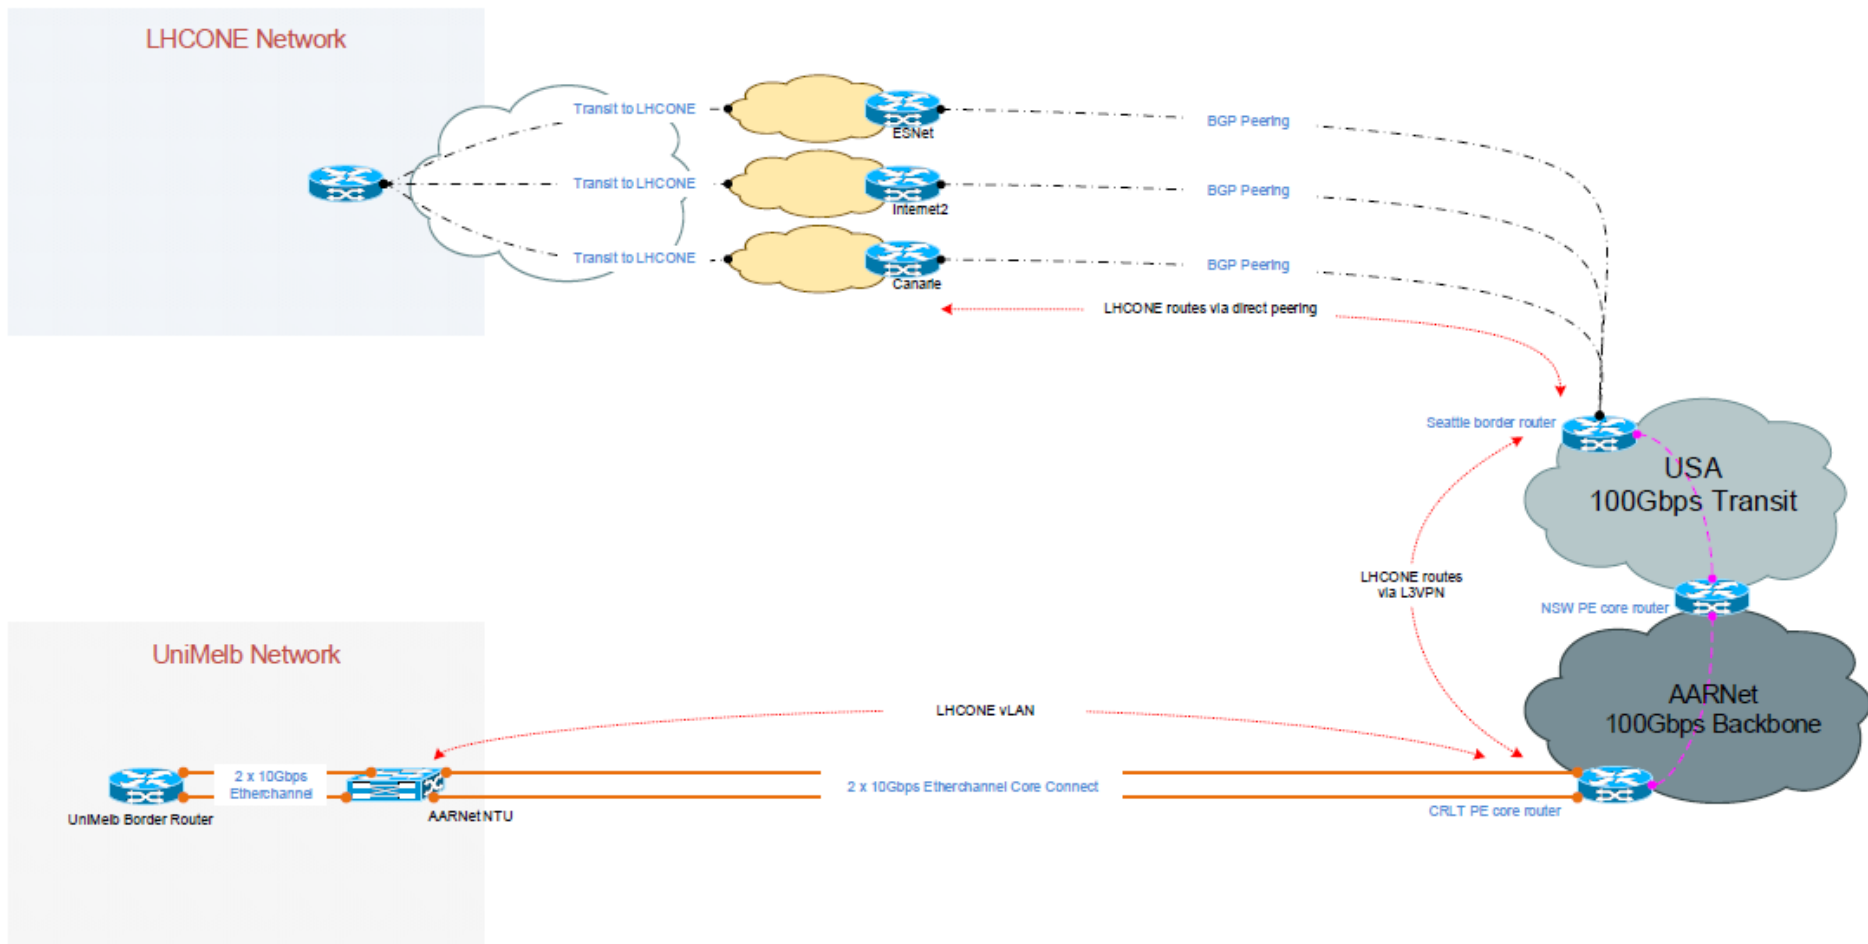
Small-ish HTCondor/ARC-CE cluster built and in production

- Running Analysis and Production ATLAS jobs
- Production Belle2 jobs

A few kinks to sort out

- How to do fairshare between Tier 3 and Tier 2 users
- How to automatically set accounting groups for Tier 3 users (easy for Tier 2)
- How to best integrate cloud (Openstack) VMs for Tier 3 and Tier 2 usage

Plan is to move our 4 Torque batch systems to HTCondor



GPU cluster

5 x Dell R730xd (3 dedicated to HEP)
2 x NVIDIA Tesla K80 GPU
1.6TB NVMe Mixed-Use Card

Usage this year:

18 users
2548924 minutes (1770 days)

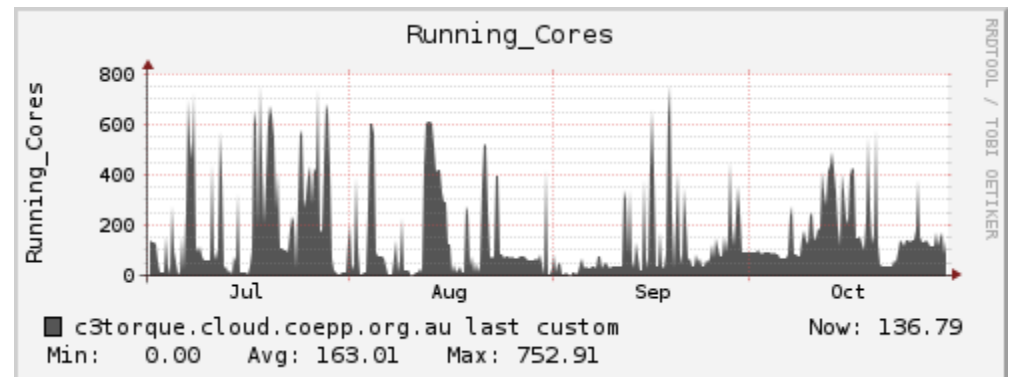
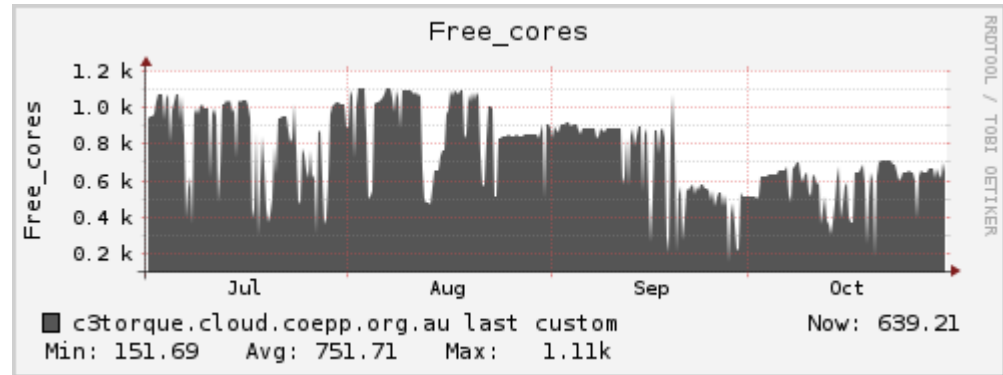
Applications:

- TensorFlow and Pytorch for rare particle identification at Belle II – recurrent networks
- N-body ($N > 200k$) dynamical simulations of dense stellar clusters
- TensorFlow and Theano for ATLAS physics analyses

Tier 3 usage

Tier 3 usage has been fairly constant with large peaks of usage throughout the year

- Backed by the NeCTAR Research Cloud
- Mainly by 4 core VMs – very difficult to start any larger VMs due to hypervisor capacity
- Mount our CephFS storage as the main storage system
- 3 cells (Melbourne, NSW and SA)



CephFS is our main data store for our local researchers

- Used to be a few NFS servers
- Ceph version 10.2.2 (Jewel)
- Has been very performant and very reliable
- Hardware is decommissioned Tier 2 storage hardware, with added SSDs for journals
- Mount it on all our Tier 3 nodes, both in Melbourne, and in Sydney, Adelaide using FUSE client
- Ceph works well over WAN, but when a network glitch happens, the FUSE client often needs to be killed and restarted

```
fsmap e41433: 1/1/1 up {0=rccephmds2=up:standby-replay}
osdmap e51280: 144 osds: 143 up, 143 in
      flags sortbitwise
pgmap v53851145: 6144 pgs, 3 pools, 82541 GB data, 40749 kobjects
      242 TB used, 176 TB / 419 TB avail
          6139 active+clean
              3 active+clean+scrubbing
                  2 active+clean+scrubbing+deep
client io 9207 kB/s wr, 0 op/s rd, 84 op/s wr
```

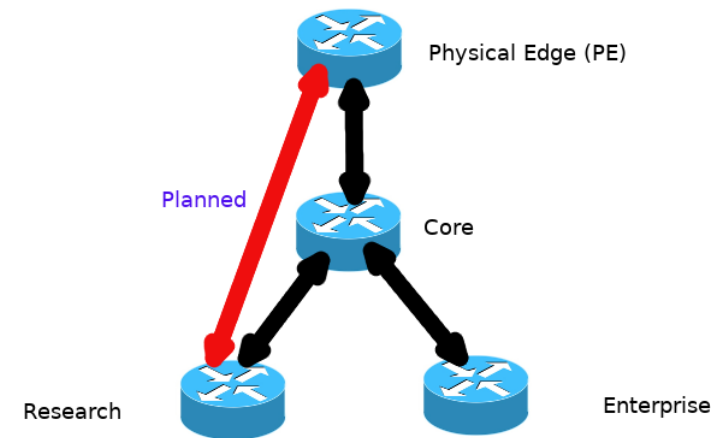
WLCG has a request to have IPv6 enabled storage for Tier-2's by the end of Run 2

University of Melbourne has had a /48 IPv6 allocation for many years now, but have never activated it on their infrastructure

Due to current network schematic, any enabling of IPv6 would require Enterprise to also have it enabled, which brings security and operational problems.

A grant request has been made to allow the research network to connect directly to the PE, and making Research a completely separate network

Once this is done, we can enable IPv6 for Research



rcteam@coepp.org.au

