CoEPP

ARC Centre of Excellence for
Particle Physics at the Terascale

Australia-ATLAS Site Report

Sean Crosby
Lucien Boland

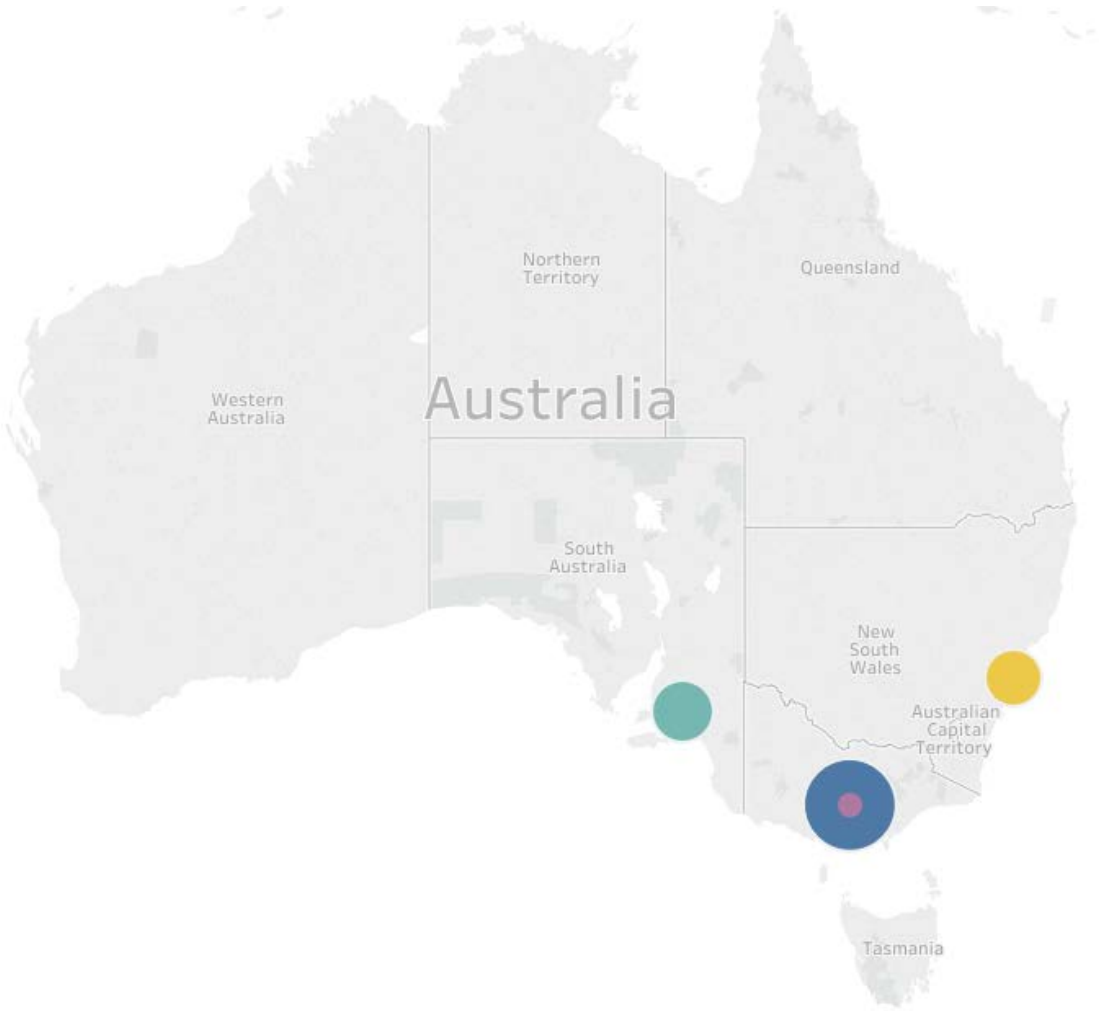
Northern Territory
Queensland
Western Australia
Australia
South Australia
New South Wales
Australian Capital Territory
Tasmania

Main heading to go here..

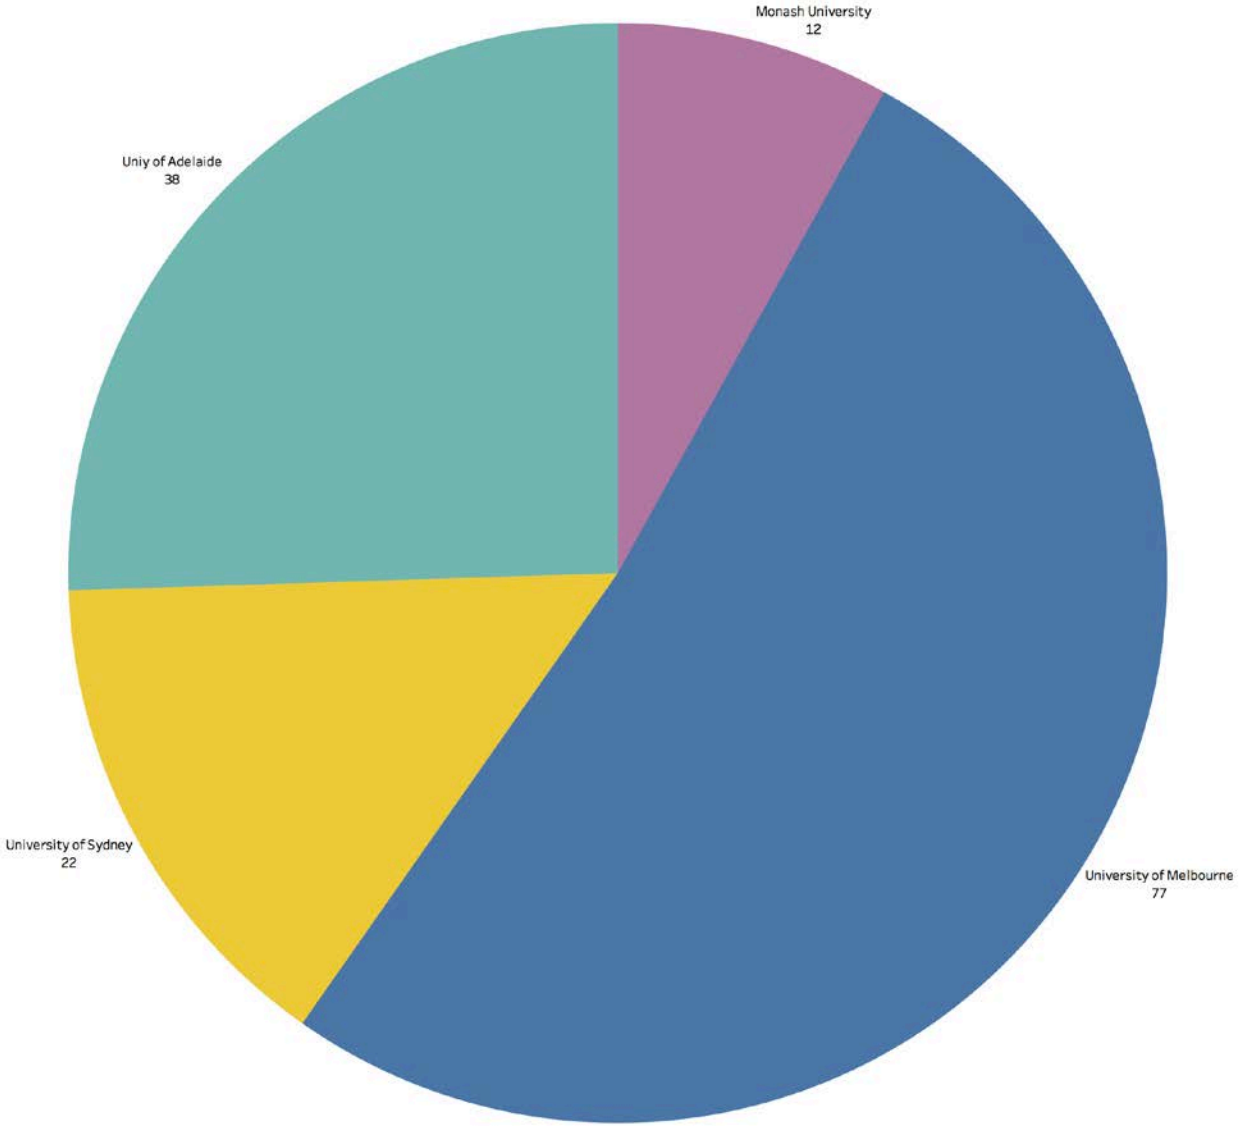

Main changes since last HEPiX

- 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(2016 HEPiX -2017 HEPiX)

Still #1 ATLAS Tier 2 site since previous HEPiX

- Improved percentages slightly
  - 99.88% Availability (2016 - 99.64%)
  - 99.93% Reliability (2016 - 99.64%)

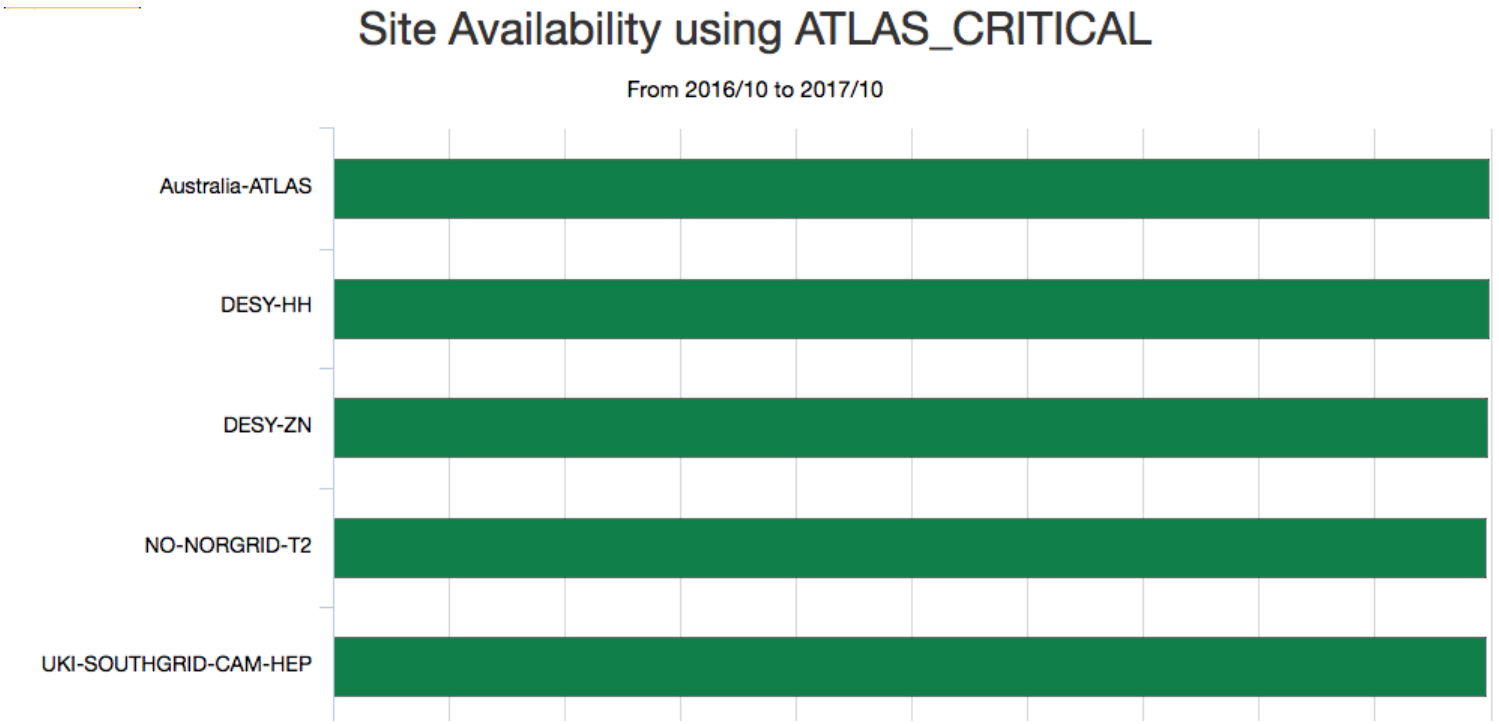

Tier 2 Growth

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

2017 and 2018 Pledge

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

HTCondor

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

- Running ATLAS test 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)
- What kind of resources (disk, memory, CPU) behind HTCondor negotiator/collector. Same for ARC-CE
- 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

LHCONE

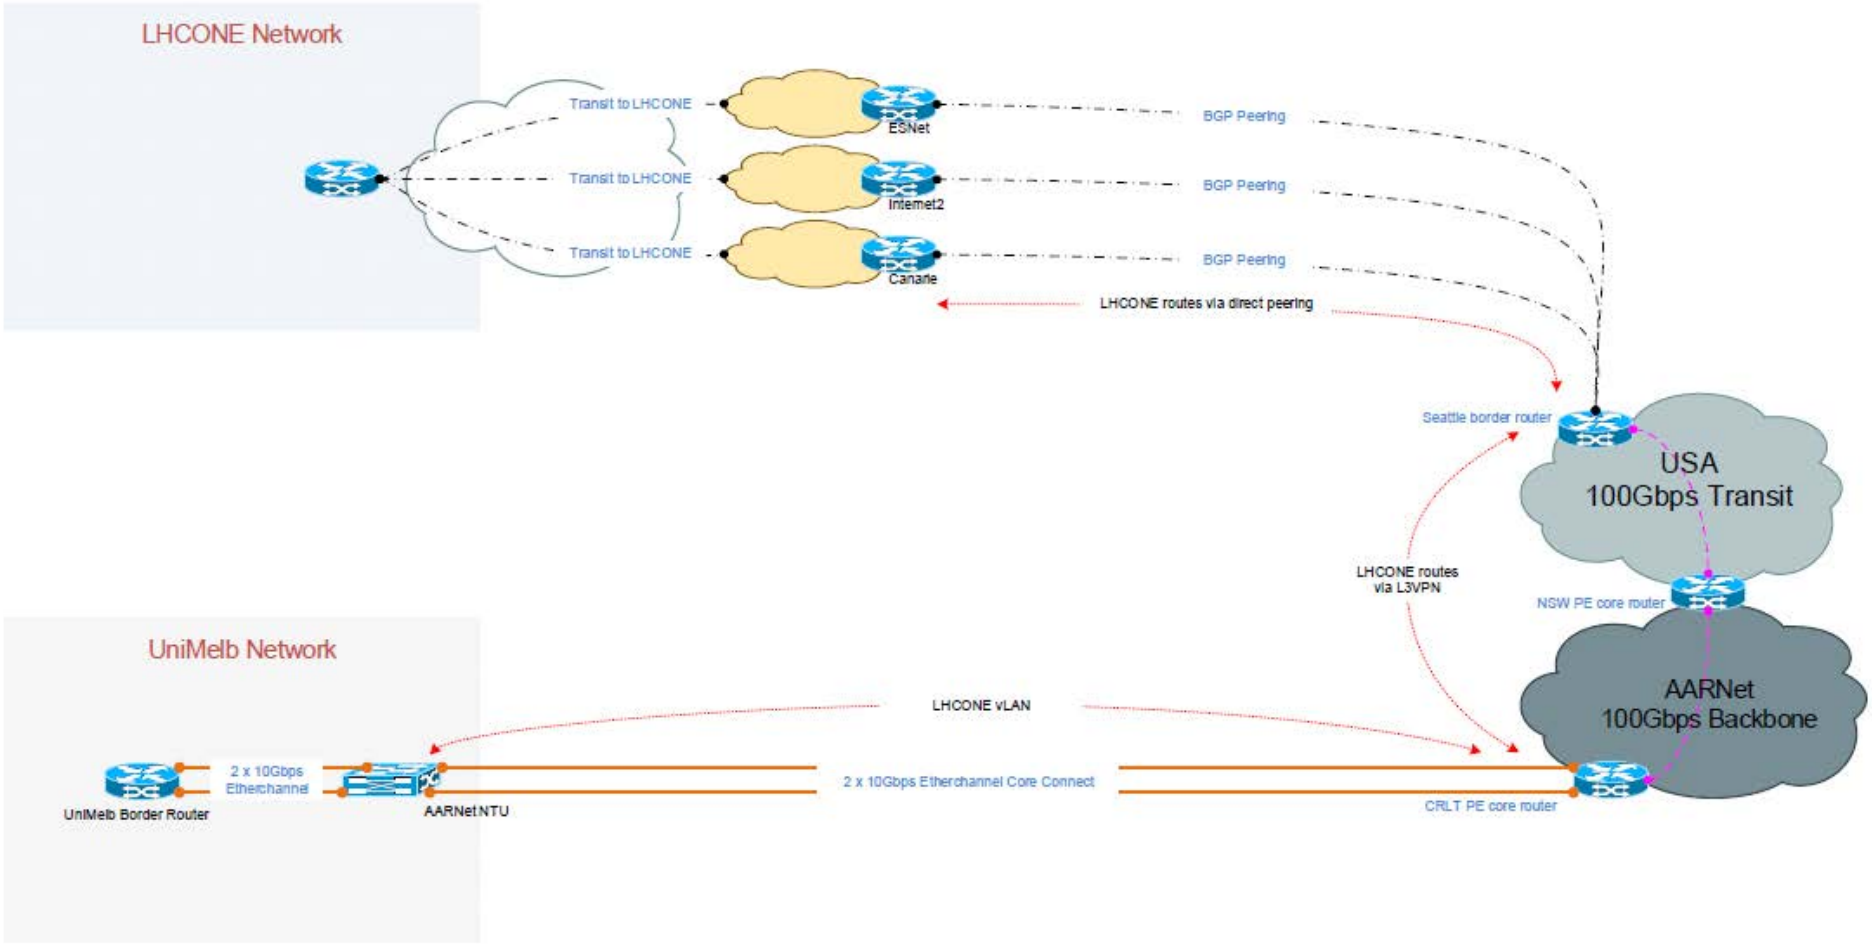
LHCONE Network
Transit to LHCONE
Transit to LHCONE
Transit to LHCONE
ESNet
Internet2
Canarie
BGP Peering
BGP Peering
BGP Peering
LHCONE routes via direct peering
Seattle border router
USA
100Gbps Transit
LHCONE routes
via L3VPN
NSW PE core router
AARNet
100Gbps Backbone
UniMelb Network
LHCONE vLAN
2 x 10Gbps
Etherchannel
UniMelb Border Router
AARNet NTU
2 x 10Gbps Etherchannel Core Connect
CRLT PE core router

Puppet 4

Existing Puppet 3 system still maintaining our SL6 boxes

Centos 7 nodes controlled by Puppet 4

Original aim was to neaten up code (had less Hiera usage than we wanted, more if statements), plus being able to “steal” other admin’s code

- 77 Puppet 4 modules used
- 57 written by us
- 20 written by other admins

GPU

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

rcteam@coepp.org.au