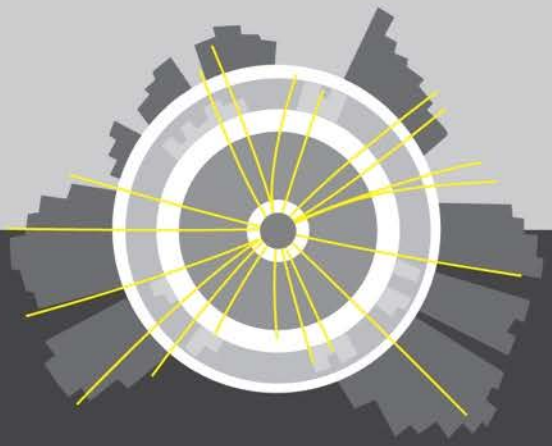

CoEPP

ARC Centre of Excellence for Particle Physics at the Terascale

Australia Site Report

Sean Crosby

DPM Workshop – 23 November 2016

- Supports ATLAS VO
- Tier 1 site is TRIUMF (Canada Cloud)
- 100Gb WAN link through Seattle
- Aim to, and delivering, 2% of ATLAS compute
- 3 full-time sysadmins
- 1200 cores
- 1.03PB storage (in main DPM pool)
- Support 100 researchers (Academics, Postdocs, PhD and Masters)

- DPM 1.8.10 (dmlite 0.7.5)/SL6
- 16 disk servers

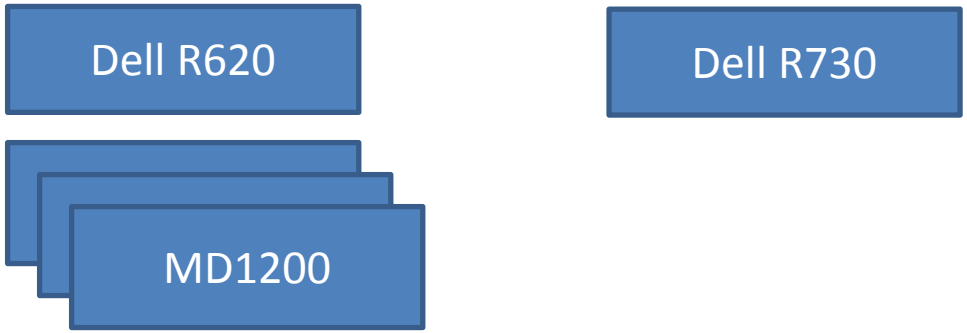

- Each drawer is a RAID-6 set
- Each FS is 10TB, made with LV's
- 2x10GbE to switch (most compute nodes 1GbE)

- Standalone head node
- 10GbE
- SRMv2.2, xrootd (FAX and local), rfio, gsiftp, webdav, dpm, dpns
- Provides storage for physical cores (on our network), and Cloud cores (see later)
- Database on dedicated DB server
  - Currently SSDs. Soon to be 10K SAS (saw very low load on DB, and old DB server out of warranty)

- 2 more DPM instances in Australia
  - B2se (Melbourne): used for Belle II compute cluster
  - coepp-dpm-01 (Adelaide): will store replica of Belle I data
- coepp-dpm-01 is backed by 200TB of NFS storage
  - WAN problems from last year have been sorted, but still a lack of performance
  - See later talking points

This year

- Still suffering Webdav crashing problems
  - About once a month
  - Browsing namespace is fine (mostly), but downloading from diskserver fails
  - Need to restart httpd for things to work again
- Changed ATLAS dump script to the updated Eygene script
  - Much slower, but much smaller /tmp space usage on DB server
- Very stable Grid site this year (again)
  - >99% A/R. 100% CPU pledge usage
- Currently signing legal documents that our NREN wants to join LHCONE
  - Should be done by end of year

- Our Centre of Excellence ends in mid-2018. We were unsuccessful in the bid for a new Centre
  - Need to look for alternate money to keep the Tier 2 maintained after 2018.
  - Part of this will involve using hardware after warranty expires
    - Currently used compute nodes are mostly out of warranty
    - Tier 2 storage is always in warranty

Out of warranty

- We have 3 spacetokens
  - ATLASDATADISK
  - ATLASSCRATCHDISK
  - ATLASLOCALGROUPDISK
- ATLASLOCALGROUPDISK does not contribute to our pledge to WLCG, but is very important for our researchers
  - Mostly data somewhere else on the Grid
  - Is it possible to contain this to out of warranty hardware

- Current use of out of warranty storage hardware is to add it to CephFS
  - Add SSDs for journals, convert RAID-6 to JBOD, then add to Ceph
  - Currently 112 drives on 11 servers. 305TB RAW (101 TB usable – 3x replicas)
  - Currently made available to Grid via STORM
- Is it possible to use this as ATLASLOCALGROUPDISK on DPM
  - Exclude DATA, SCRATCH disks from being on it
  - What about with new, SRM-less , DOME-ful world?
  - Is it possible to add multiple Gridftp/xrootd/Webdav servers to service this one filesystem?
    - Would also help with our NFS situation in Adelaide

- Logging in 1.8.10 is quite bad for dmlite components
  - Looking forward to 1.9.1 to see the improvements
- We currently take database dumps every 12 hours
  - Currently 1.5TB of database archives going back to 1/Jan 2015
  - We had to use the database dumps (and SRM log files) to try to find out why directories in SE namespace had wrong permissions, so they sometimes can be helpful
  - TRIUMF just used their dCache billing logs from 2015 to see why a file ATLAS says was on their SE really wasn't
  - Is dCache-like billing logs something of interest to other sites, and feasible to implement?

scrosby@unimelb.edu.au