



COEPP

ARC Centre of Excellence for
Particle Physics at the Terascale

Australia Site Report

Sean Crosby
DPM Workshop - 9 October 2014

- Supports ATLAS VO
- Tier 1 site is TRIUMF (Canada Cloud)
- 10Gb WAN link through Seattle
- Aim to, and delivering, 2% of ATLAS compute
- 3 full-time sysadmins
- 1000 cores
- 920TB storage (in main DPM pool)
- Support 100 researchers (Academics, Postdocs, PhD and Masters)

- DPM 1.8.8/SL6
- 9 disk servers
- Each FS is 10TB
- 2x10GbE to switch (most compute nodes 1GbE)

- Standalone head node
- 10GbE
- SRMv2.2, xrootd, rfiio, gsiftp, webdav, dpm, dpns
- Provides storage for physical cores (on our network), and Cloud cores (see later)
- Database on dedicated DB server
- Everything but the DPM config is puppet controlled (still using YAIM because of dependencies)

- Now part of ATLAS FAX
 - apart from bug in N2N plugin, working great
 - no noticeable change to load on disk servers
- Added two test beds (serving real data) – more on it later
- Educated our users to use xrootd when possible in their Tier 3 code

- Following on from last year, our cluster is still heavily unbalanced
 - Every year we drain the oldest disk servers, and the data drains nearly directly to the newest
 - Nearly all writes heading for the 2 newest disk servers
- Large amount of dark data remains
 - waiting on Rucio dumps and tools to identify where the data is

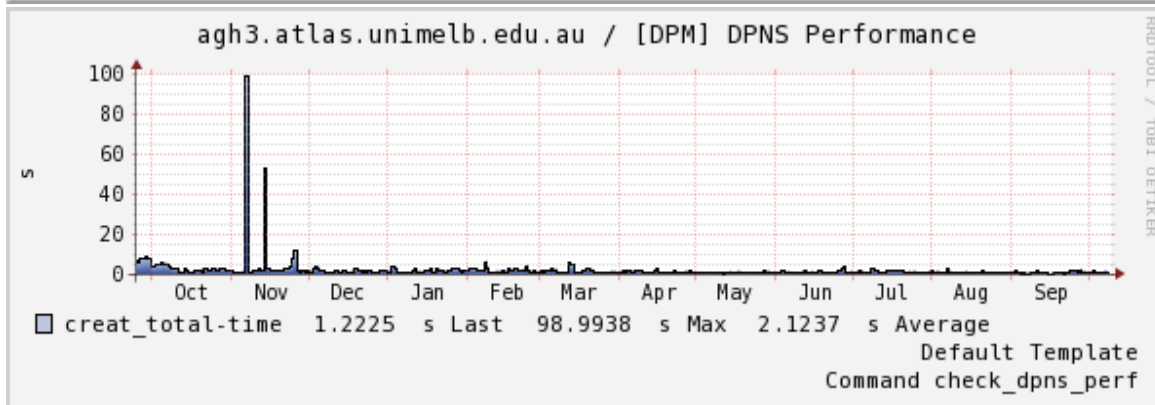

- We are part of Webdav federation with all other Canadian sites

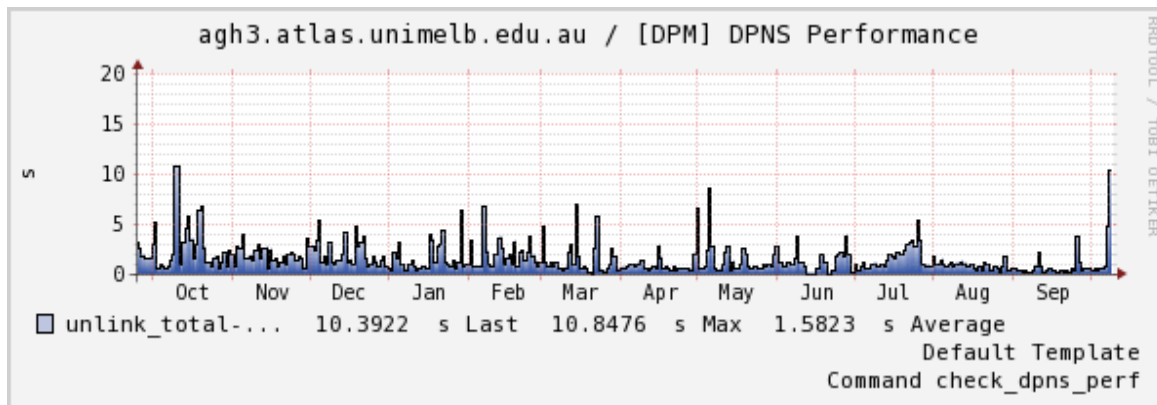



- Federation running using UGR (based on dmlite) – Fabrizio will talk more on it
- Current purpose is to use it for data inputs for Cloud computing sites – UGR has a GeoIP option to select closest replica. 3 ATLAS cloud sites in CA.
- A few problems with Australia
 - Our SE falls out of the Federation regularly due to unknown reasons
 - We also see SSL connect errors when connecting remotely (fine locally)
 - More diagnosis needed

- Also need to fix ATLAS pilot to work with Federation – should be simple enough
- Great work by Ryan Taylor (UVIC) and Fabrizio in helping to identify new features we'd like in the UGR

- We now have a full year of monitoring of DPNS, DPM, GridFTP performance, using the lcgdm-nagios plugins
- Only a couple of issues caused by faulty disk server





- Our 3rd sysadmin starts in November
 - Currently I am only sysadmin who knows DPM
 - Making docs in our internal wiki on common faults and error logs, and solutions
 - e.g.
 - grep "Restarting query after lock wait timeout" /var/log/dpm/log
 - grep "Host not known" /var/log/dpm/log
 - grep -i mysql /var/log/d*/log
 - But, don't be afraid of
 - grep SRM_FAILURE /var/log/srmv2.2/log
 - [SE][GetSpaceTokens][SRM_INVALID_REQUEST]
http://agh3.atlas.unimelb.edu.au:8446/srm/managerv2:dpm_getspacetoken: Unknown user space token description
 - Actually a permission denied on space token
- Is it worth having it on main DPM website?

- 2 more DPM instances in Australia
 - b2se: used for Belle II compute cluster
 - coepp-dpm-01: will store replica of Belle I data
- coepp-dpm-01 is backed by 200TB of NFS storage
 - So far so good
 - Having problems with WAN connection, so performance testing not yet completed

- Given 500TB of storage from federal government
 - wanted it accessible to Tier 3 users like regular space, as well as accepting transfers from Grid
 - running Ceph with CephFS
 - STORM provides Grid functionality
 - a few performance problems, but otherwise great



scrosby@unimelb.edu.au

