

Lancaster: The more things change...

HEPSYSMAN SITE REPORT

Matt Doidge, Peter Love
24th June 2019



- Grid Tier 2 that utilises the University shared cluster (HEC), and now an IRIS site.
- Matt holds the Tier 2 fort, Peter keeps the local users in check. Our counterpart Mike Pacey stops us taking over the HEC completely as our counterpart.
 - Robin has left us to fight bad code in Manchester.
- EPP group primarily involved with ATLAS and T2K, and the grid site reflects this. Time a changing though - DUNE, MICROBOONE and others are climbing the ranks, and we see astronomers doing stuff on our cluster thanks to IRIS (SKA, LSST, CTA).

- The EPP cluster is fairly typical.
 - SL desktops, HTCondor batch system, NFS (well exported ZFS) shared areas, a web server, a beefy machine to log into and run stuff on.
- Because users always want more space Peter is bringing online a new ZFS scratch space.
 - Has anyone had any luck using ZFS quotaing to keep users under control?

What's new at the Tier 2 Zoo?

- Still using a CREAM CE, DPM SE, cobbler/ansible to deploy stuff when we're good. (Almost) Everything that matters is C7.
- A change in vendors means new compute has a different badge on the front and we needed to learn racadm.
 - Storage nodes have changed a bit more, with fewer drives in smaller chassis (and M2 system disks - which worked out of the box).
- Storage is totally v6'd, we're slow with the other services.
- DPM is 1.12, DOME-flavour (that was fun...) and running on the University VMWare infrastructure - which is working well.
- Monitoring is still using icinga and ganglia, our dashing dashboard died. Peter is replacing with a solution using grafana and statsd.

- Lancaster is now an IRIS, err, site? An IRIS place where there's hardware.
- But we're lucky, our IRIS kit is more or less just an annex of our gridsite.
 - Our IRIS hardware is to be presented to the world as simply more grid enabled kit, just with more priorities to different VOs (sometimes from somewhat different communities)
 - So far our main takers are LSST, SKA, LZ, DUNE and CTA
 - On our SE they've been given a separate DPM "Pool".
 - Eventually we want to spin off the IRIS stuff with a separate queue and CE.

- I had a ambitious plan to come here with a talk on our HTCondorCE progress.
 - But the Universe has had other ideas.
- But my initial experiences, guided by the mighty Ste Jones, have been positive.
- Why HTCondor?
 - Because I never had a feel for ARC.
 - Because `condor_ce_submit` is a thing.
 - Because the Ste and friends have slain most of the Bugbears (BDII, APEL).

- The UI/WN tarballs are still being looked after by me.
- (Although I've sneakily EOL'd the SL6 versions.)
- But outside of /cvmfs/grid.cern.ch they're homeless - we never found a place to keep them online after the loss of the EGI repo.
- Any requests for features?
 - I recently tried to put in the condor_ce clients, but they're full of absolute paths which break any relocatable distribution.
 - I'd like to reduce down to 1 tarball in the future - a combined UI/WN tarball.

The Melty Man Cometh...



Call a Sparky!

- We've had (avoidable) electrical problems crop up as our machine room nears its first decade in service.
- Loose connections like the above meant we had to replace all our commando sockets.
- Similarly overly bunched cables under the floors have caused unwanted heating.
 - The fix is to replace the cables with thicker versions.
 - This work is ongoing live alongside a PDU replacement exercise - steadier hands than mine are swapping out cables as we speak.

Any questions?



Matt Doidge, Peter Love
24th June 2019

