

Lancaster HEP Computing

HEPSYSMAN 2010

Matt Doidge, Peter Love & Rob Henderson.

A familiar story...

- In our offices you'll find a bunch of SL5 linux desk tops.
 - NFS mounting user's home areas.
 - Which are of course backed up nightly.
 - As well as several TB of user disk space.
 - Which gets gobbled up pretty darn quick.
 - All are part of the local batch system.
 - This batch system is supplemented by a handful of decommissioned grid nodes.
- Users are also provided access to a webserver, a video conferencing suite and access to not one but two Grid Uis.
- In the basement you'll find one of those Tier-2 Grid thingies.

Recycling isn't just for beer cans.

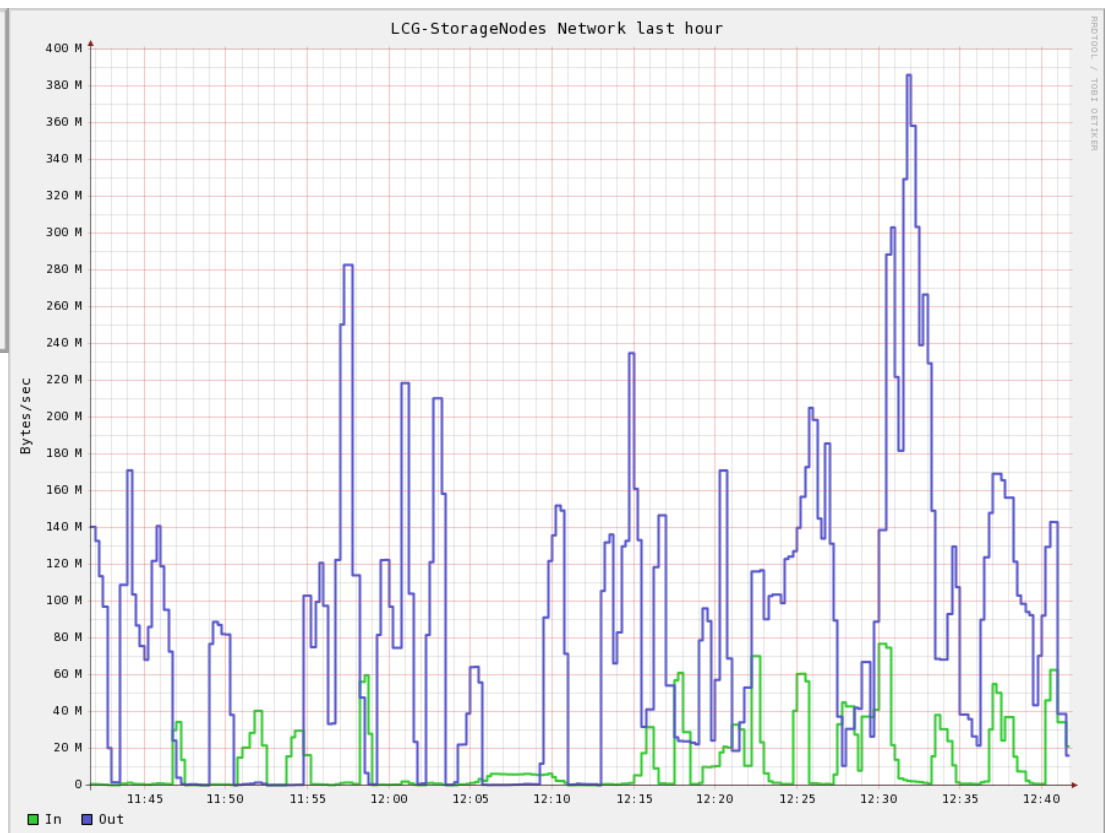

- Decommissioned Grid Storage retasked as user scratch space.
- Along with the racks they came in.
- Previous success retasking Grid test worker nodes to suppliment local batch system.
- Local hardware purchased by piggy-backing off of larger Grid orders in the past, possibility for future financial collaborations.

Lancaster's Tier-2 Grid Thingie.


- Part of Northgrid.
- Computing Element running torque/maui.
 - Just shy of 600 jobslots.
 - Most jobslots come from dual and quad chassis Wns.
 - Workers configured using an nfs mounted tarball.
- Storage Element running DPM.
 - 300 TB online, 70 TB “in reserve”
 - 24-bay supermicro disk pools with 22x 1- or 2-TB disks.
- 10 Gb LAN, 1Gb light path to RAL.
- Everything held together by sweat, tears and cfengine, backed up with Bacula.
- Move to Virtual Machines for low load services.

Keeping an eye on it all.

- We use Ganglia, Nagios, syslog and homegrown (or brazenly stolen) scripts to keep tabs on everything.



Storage used by group



- alice_Role_lcgadmi
- atlas (1.60%)
- atlas_Role_lcgadmi
- atlas_Role_pilot (
- atlas_Role_product
- atlas_it (0.00%)
- atlas_uk (0.00%)
- atlas_uk_Role_powe
- biomed (0.13%)
- dteam (0.00%)

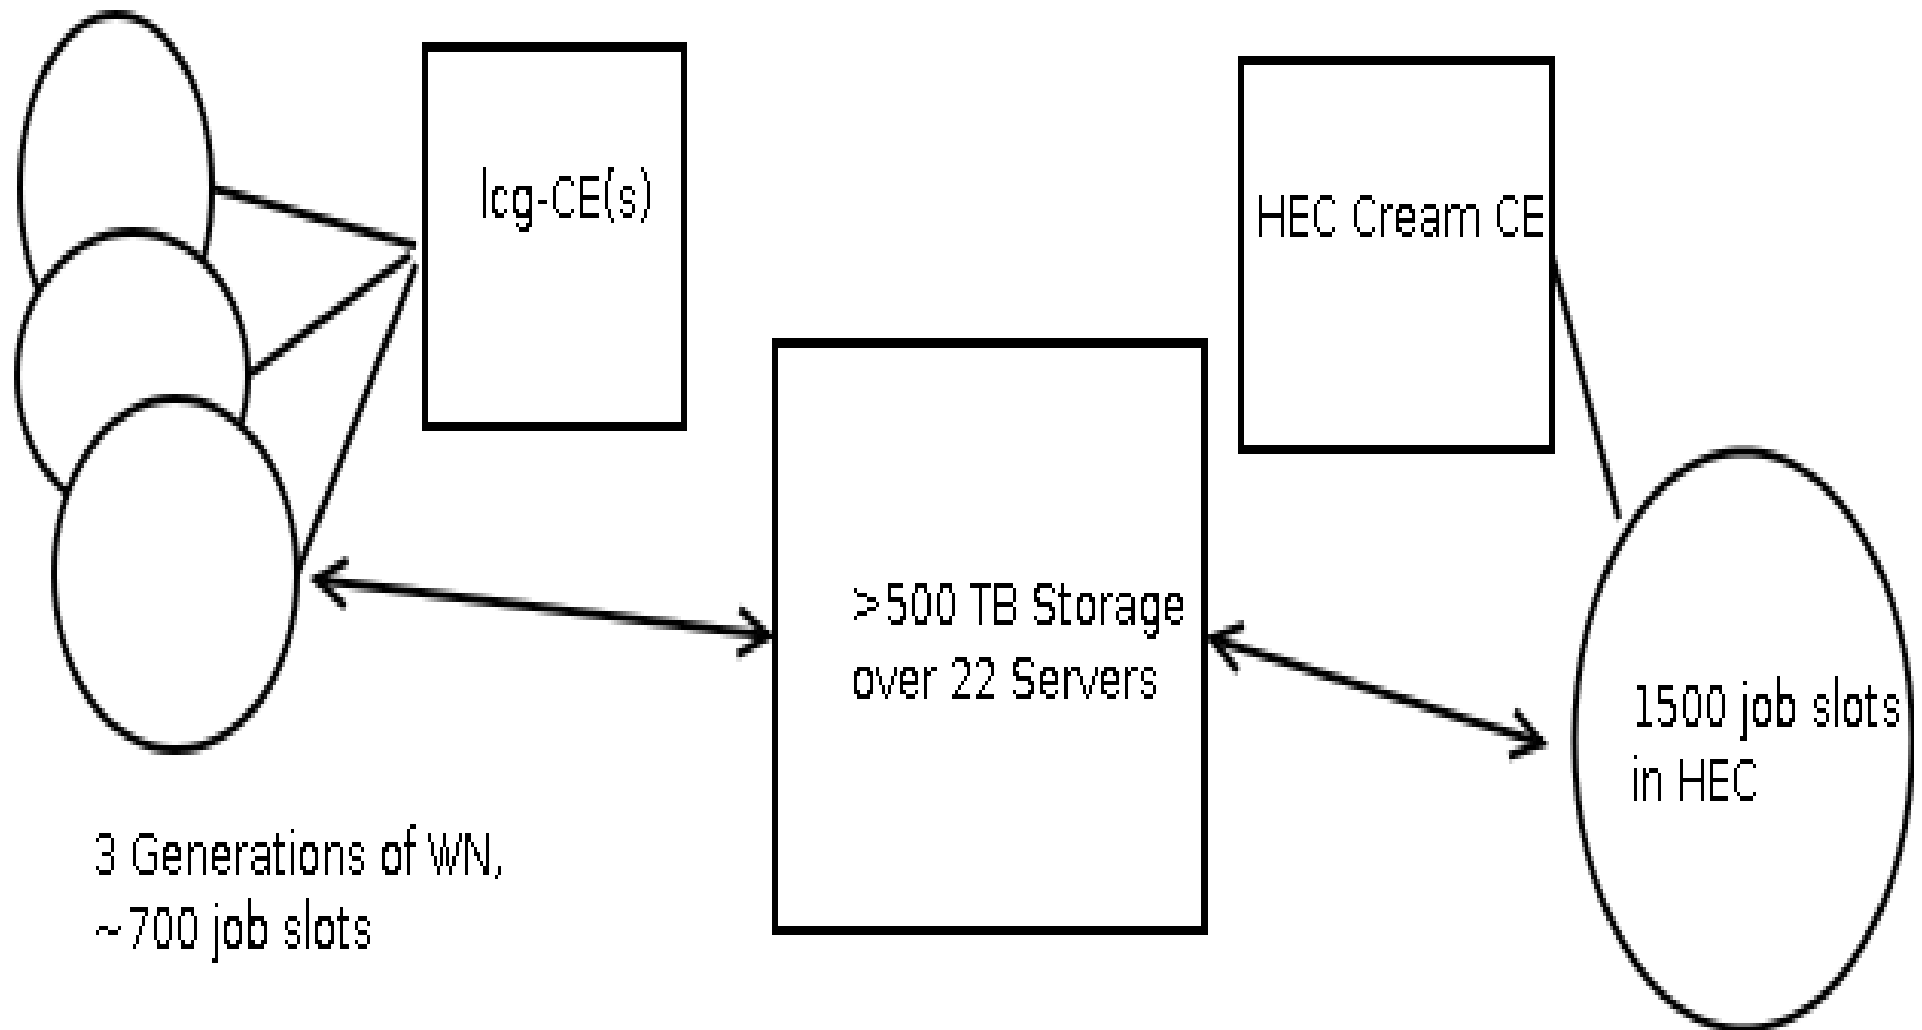
Highway to HEC.



All hail Roger, Lord of HEC.

- Lancaster will soon(ish) open it's HEC facility, a joint venture between HPC interests at the University.
- The Tier-2 is a major player in this, with our indomitable boss Roger appointed director.
- It will give us access to:
 - (a share of) 1500 job slots in a LSF batch system.
 - A Panasas storage shelf to serve tarballs and software.
 - 245 TB of extra storage all for ourselves
 - A proper machine room!

After HEC we'll look something like...



Lancaster's Grid Gotchas.

- Contention between lightpath traffic and other WAN traffic, causing connection troubles.
- Performance issues trying to serve tarball and experiment software from the same box.
- Grid backups are getting huge – 1TB for configs and 6 months of logs – not including the DPM database backups.
- Data access patterns by jobs to our SE need to be understood so we can tune access methods.

The Never Ending To-Do List.

- Beefing up our Janet link.
- Streamlining batch system access for local users.
- Rollout new services (we're still a CREAMless site).
- Understand and tune data access patterns.
 - We'll need to justify our fairshare.
- We still have to commission all that shiny new hardware when it arrives...
- **And we've probably got 50 bazillion other things to be getting on with.**