

Lancaster Site Report

Matt Doidge & Peter Love.

HepSysMan June 30th 2009

Our Kit in a Nutshell 1: The muscle.

The Old:

- 200 Dual-core Xeons.
 - 400 cores.
- 50 TB in SCSI mounted Disk.

The New:

- 16 Dual Quad-core Supermicro twins.
 - 256 cores.
- 120 TB in 6 24-bay raid boxes.

Our Kit in a Nutshell 2: The Brains

- 2 CEs, one for the old cluster, one for the new.
- DPM Head Node for the SE.
- 2 “Grid-run” UIs.
- A multi-tasking Mon-Box (Ganglia, Nagios, Site-BDII, Accounting DB, Thermostat, syslog).
- 2 NAT boxes, one for each cluster.
- SSH Gateway/workbox.

Our Kit in A Nutshell 3: The Central Nervous System.

- Fabric management: cfengine
- Installation: PXE/Kickstart
- Backups: Bacula
- IPMI on the new kit.
- Remote power management of Old Kit.
- External Network on own subnet.
- Internal Network Class A.
- Main switches are HP Procurve, planned and administered by ISS.

Grid Plans 1-What's better than twins?

- 2009 Kit Purchase (Provisional).
 - 200-300 TB of new Disk- but may need 2TB HDD to make this happen.
 - 10 GigE Infrastructure.
 - 256 Nehalem cores with 3GB RAM/core, perhaps in Supermicro Twin^2 chassis.
 - Replacement servers for SE/mySQL – finally a split DPM headnode!
 - 16 core “virtual host”.

Grid Plans 2 – A new home.

- Next year we'll be moving to a new home.
- A proper made-for-purpose machine room.
- Eventual consolidation of our primary resources.
- Water cooled racks.



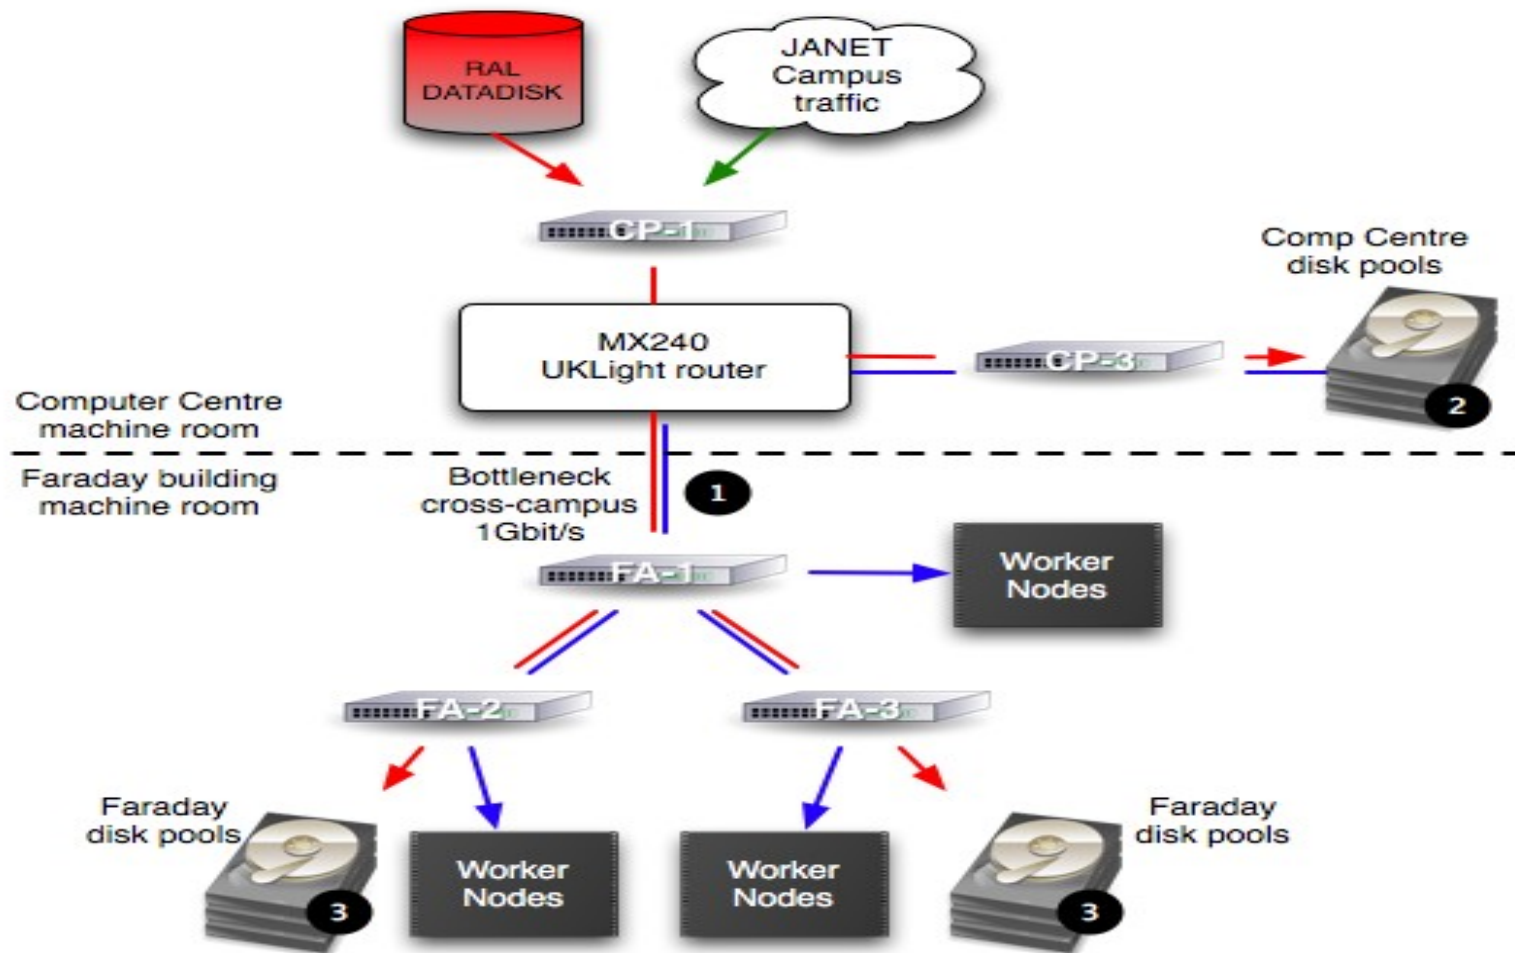
Grid Plans 3 – New friends.

- We'll be moving to our new machine room as part of the University's new “HEC” group.
- The HEC is (will be) the amalgamation of future GridPP and Uni HPC resources, as well as additional support from User groups.
 - University recognition for our GridPP work.
 - Future access to greater resources, with a voice on the steering group for the University's central HPC resources.

Grid Plans 4: Tier 3 (2.5?).

- Some current “Tier-3 like” activities
 - Grid brought and run 20TB NFS server.
 - Grid-run UIs.
 - Recycled Grid nodes for local batch system.
 - But we need to do more-
 - Local submission to CEs
 - WN mounted scratch area.
 - Posix-like access to SE from Desktops (Global File-system?).

Step '09 Experiences 1: Spot the Bottleneck(s).



Step '09 Experiences 2: Going Global.

- Atlas experiences during Step '09 showed great improvements in Job efficiencies when using file:// type data access.
- Local users would also like file:// type access.
- So we would to consolidate our storage pools into a single global files-system.
 - But who, what and how? Those running global file-systems can expect to have their brains picked this week.

Other Trials and Tribblelations: Fabric.

- We had the AC upgraded in our main machine room.
- But we're "living on the edge" w.r.t power.
- We've obtained the rights to two new water cooled racks for the new kit, but in yet another room.
- Older WN 3-Com switches keep dying, knocking out chunks of old WN at a time.

Trials and Tribblelations 2: The Amazing Ever-Growing To-Do List.

- LDAP for users.
- Even more monitoring, with some automation-but not too much!
- Exploring the use of Virtualisation-but again not too much!
- Culling of old equipment.
- Alternatives to Torque/Maui?

Trials and Tribblelations 3: Recent Fun and Games.

- We spent a few months not publishing the accounting on our new CE due to a misconfiguration. The Moral of this story is check your accounting regularly.
- The delay in DPM 1.7/SL5 compatibility bit us on the backside.
- DPM in general is proving a bit too “feature poor”, seen when dealing with draining and hot files.