

Lancaster Site Update

Matt Doidge.

Contributions from Alex Finch, Peter Love and Rob
Henderson.

HepSysMan June 2008
RAL

Lancaster- people

- Matt Doidge & Peter Love handle the Grid side of things.
- Rob Henderson and Alex Finch handle the local machines (and users).
- Brian Davies contributed a lot to setting up the grid network structure.

Lancaster HEP Machines

- 4 Servers:
 - Web/Mysql
 - Nis/Mail
 - Printer/Batch
 - Disk/Tape
 - + tape robot
- ~13 Workstations
 - SL4
 - Run batch jobs as well as act as desktop.
 - Varying specs and makes.

Other Local HEP services

- Zydacron video conferencing suite.
- Polycom conference phone.
- Polycom Laptop.
- 4 printers.
- 8-node batch farm from old grid test rack.

Lancaster Grid Overview

- Computing Element
 - 200 dual-core worker nodes in 4 and a half racks.
 - Xeon processors with 2 GB DDR RAM.
 - 80GB SATA drives.
 - 400 kSpecInt
 - CE itself is one of the old gridpp boxes!
- Storage Element
 - ~80 TB storage space.
 - Migrating from dcache to dpm.
 - 8 Disk-servers with 10-TB each.
 - All storage direct attached via SCSI.
 - Old but reliable dpm headnode.

Other Grid Machines.

- Monbox and BDII, also runs Nagios and Ganglia monitoring of our nodes.
- UI for local users.
- Gridmon box.
- SSH Gateway.
- Backup Server- Bacula based archiving of important information.
- Fabric Box- cfengine, subversion repository of configs, NAT, YUM repository mirroring.

HEP Cluster Problems

- Atlas batch jobs were eating desktop resources
 - Upgrade of all desktop machines RAM.
- Evo refuses to work with the Zydocron videoconferencing suite.
 - Issue unresolved, despite collaborating with the EVO developers.
- All nodes to be upgraded to SL4.4 within the near future.

Grid Problems

- Running out of Space, Power and (most importantly) Cooling.
- Running hot seems to promote harddrive failures.
 - 2 downtimes this year due to the system disk popping on a server.
 - Estimate a rate of 1 WN a month blowing its disk.
- On top of this the warranty on our kit has run out!
- Due to kit age replacements parts are getting harder to find.

But things a looking up

- Space problems have been alleviated by shipping old kit to Kenya.
- We're getting the electricians in to provide more power.
- And hopefully we'll be able to nab another chiller from next door
 - Making us ready for new purchases next month.
- And anyway next year we're moving to a bigger pad.

New Grid Purchases

- Within the next month we hope to purchase extra storage and CPU.
 - 80-160 TB of storage and 16-32 high-RAM dual quad-core CPUs, numbers dependant on price.
- Also looking at purchasing replacement core service boxes- new CE and DPM headnodes with higher spec and greater resiliance (mirrored system disks!).

A change of SRM

- Migrating from DCache to DPM.
 - DCache provides more features than we need.
 - DPM a much simpler solution.
- DPM been in production for a few months now, no real problems.
- Continued process of data migration and disk-server “cannabilisation” over the last few months.
- DCache to be finally decommissioned at the end of this month.

A New Home

- Lancaster University are building a new ISS facility.
- Within this facility we get a 48-rack server room!
- Unknown if this is going to be a water or air-cooled facility at this time.
- But we should be able to move in during the Summer 2009.

Tier 2 and Tier 3 working in harmony...?

- Increased collaboration between the grid and local clusters.
 - “Grid supplied” UI.
 - Grid Test Cluster reused as a batch cluster.
 - Local user submission as a possibility for a future cluster.
 - Sharing of machine room space in the future, and working together on purchases.

In conclusion

- Business as usual at Lancaster.