Cambridge HEP Group - site report

May 2006

John Hill

UNIVERSITY OF CAMBRIDGE

Group hardware (non-GRID)

- Have a total of five DELL file servers:
  - 3 existing servers running SLC306
  - 2 brand new servers (will be) running SLC43 (from SRIF)
  - total of ~30TB disk space served
    - 20TB SATA-FC RAID (also from SRIF - still being commissioned)
    - 5.8TB SCSI-SCSI RAID
    - 3TB IDE-SCSI RAID
    - Rest JBODs
  - LTO-3 autoloader for main archiving
  - Legacy DLT8000 autoloader
  - all UPS-protected

Group hardware (non-GRID)

- RAID experiences:
  - IDE-SCSI: reported last year the serious problems with these arrays – things have improved recently, mainly because the usage level is not so high!
  - SCSI-SCSI: once the weak disks had failed (~5 in first 9 months) we've had no problems.
  - SATA-FC: ask me next year…

Group hardware (non-GRID)

- 67 desktop PCs (give-or-take...)
  - 39 Linux (SLC306)
  - 27 Windows XP/2000
  - 1 Windows 95 (for our probestation)
  - wide range of performance (400MHz Pentium II up to 3.4 GHz Pentium 4)
- 1 MAC desktop to keep me on my toes.
- 33 PC and 10 MAC laptops registered (but certainly not all are real machines...). Purchased with a confusing mix of group, college and private funds. PC laptops run a mix of Linux, Windows (and frequently both).
- Also a few PDAs in use (though not in the numbers I was expecting a year ago).

Group Hardware (non-GRID)

- Printers – unchanged from last year…
  - 2 LexMark B&W laser printers (20ppm+) – 1 supports A3 for (e.g.) CAD drawings
  - HP Color LaserJet 4550 (for paper) – needs replacing soon.
  - 2 HP Business DeskJets (for transparencies)
  - Epson A2 colour inkjet (specialised printing: CAD, posters,...)

Hardware for GRID work

- Long-standing: 20-CPU analysis farm (MIMCluster from Workstations UK)
  - 1.13 GHz Pentium III CPUs
  - Linux SL3
- Added a further 10 systems a year or so ago:
  - 2.8 GHz Xeon dual-CPU DELL PowerEdge 1850 servers
  - Linux SLC3
- Also 4 systems provided by GridPP at about the same time:
  - 2.8GHz Xeon dual-CPU Streamline Computing servers running SL3
  - Front-end boxes for LCG2 deployment.

- 3TB IDE-SCSI RAID specifically for GRID work (i.e. in addition to that mentioned earlier).

- Also have 4-node WXP cluster to support Microsoft-funded student in Windows-based GRID work.

CamGRID

- Collaborative project involving (mainly):
  - Cambridge eScience Centre (CeSC)
  - National Institute for Environmental eScience (NIEes)
  - eMinerals project
  - University Computing Service (UCS)
  - Cambridge HEP Group

- Uses Condor as batch system, just as LCG farm and Cambridge HEP cluster do. Status can be seen at

  http://www.escience.cam.ac.uk/services/camgrid-status.cgi
- Cambridge HEP cluster can already flock to CamGRID, though use of this is currently limited due to differences in the environment.
- Plan is eventually to have the LCG farm able to flock to CamGRID, to allow us to increase our provision to LCG.

- CamGRID were successful with a SRIF bid – a tender has just started to procure the kit. ~128 CPUs will probably come to HEP under this.

Network

- Wired:
  - departmental network based mainly on switches from Extreme Networks
  - Gigabit Ethernet fibre-optic backbone between switches
  - a minimum of Fast Ethernet to all desktops
  - Gigabit UTP to recent fileservers (using Netgear gigabit switches)
  - Gigabit connection from department to campus backbone
  - 10Gbps Ethernet on campus backbone
  - Campus network connects to EastNET via a ~12.5Gbps link and hence to SuperJANET.
  - departmental network in process of moving to VLANs (this is rather slower than intended due to the existing IP allocations being rather "random" across a number of groups).

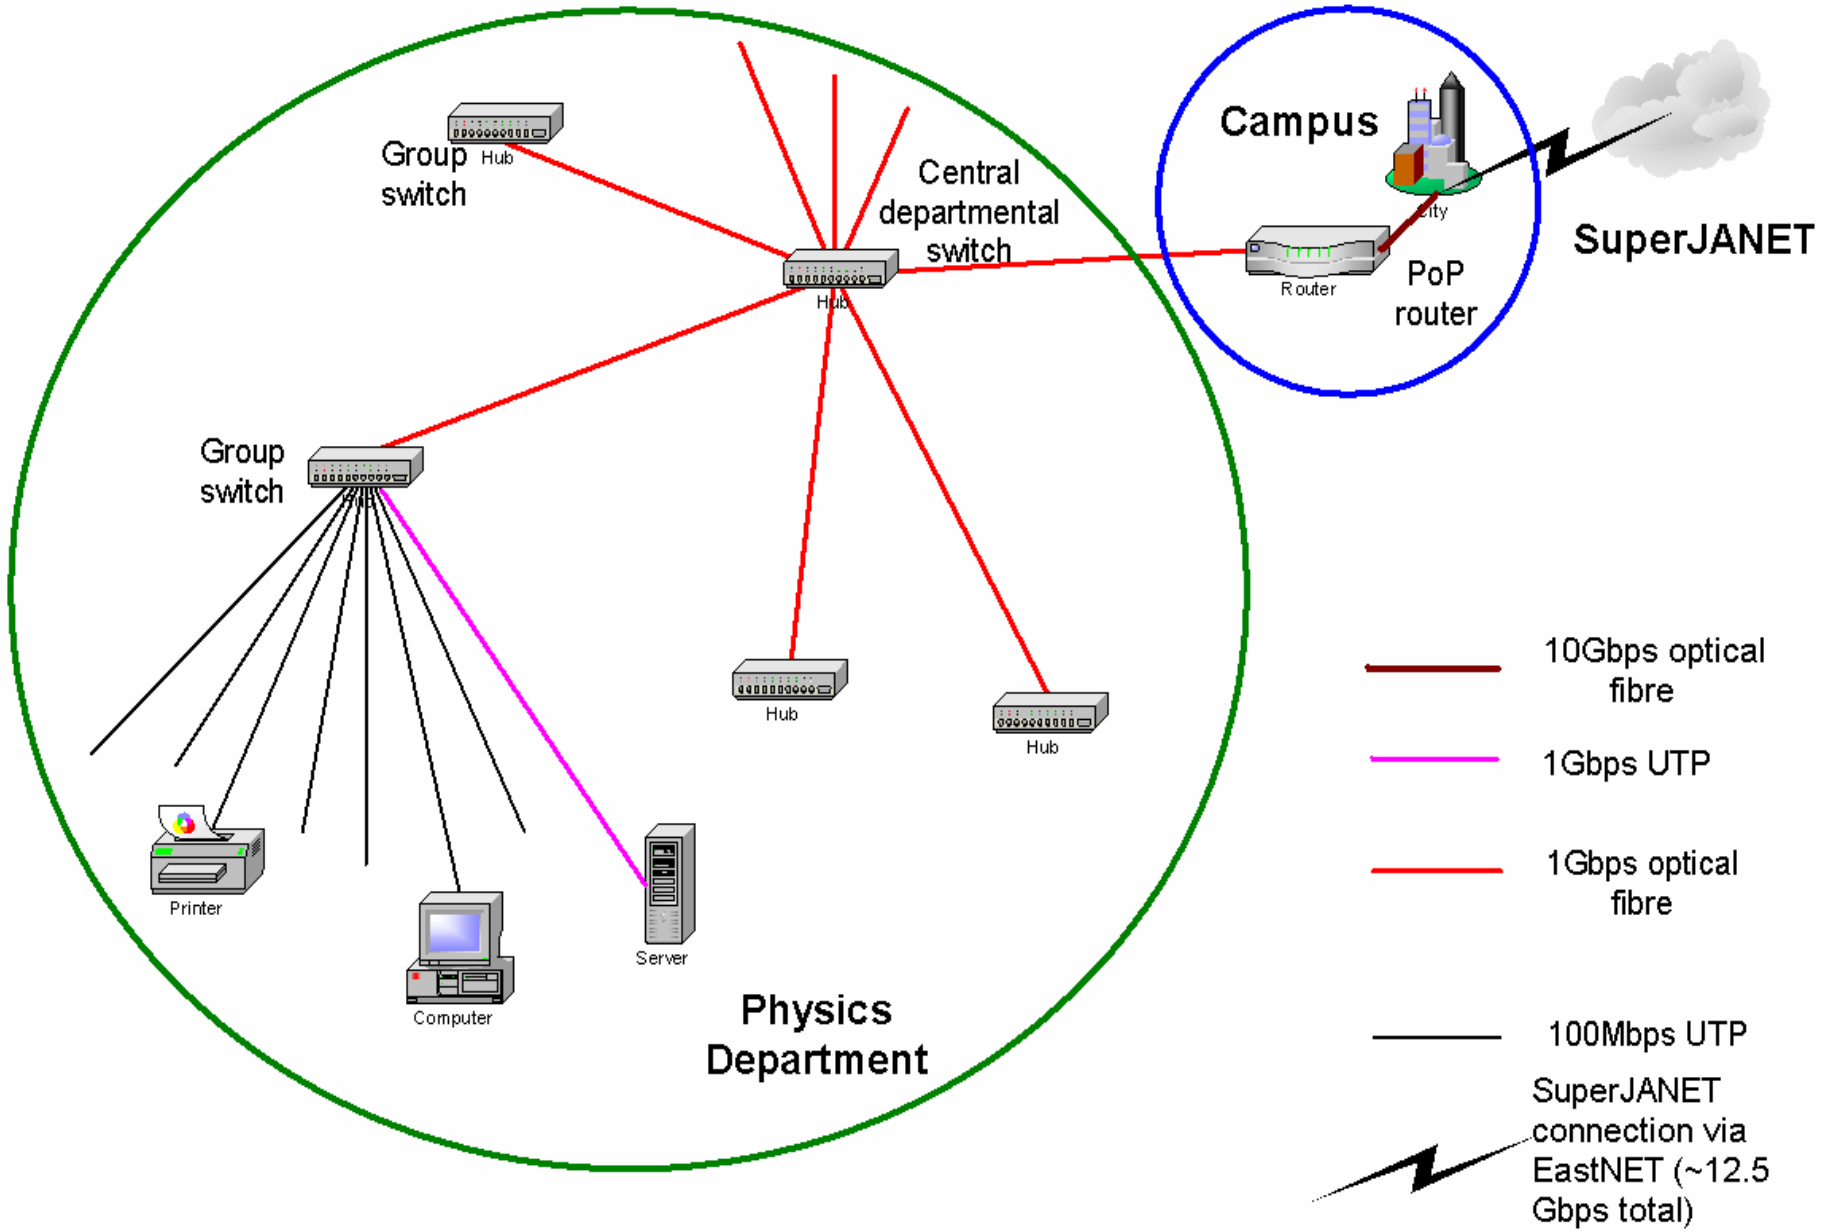
Group switch
Hub
Central departmental switch
Hub
Campus
City
SuperJANET
Router
PoP router
Group switch
Hub
Hub
Printer
Computer
Server
Physics Department
10Gbps optical fibre
1Gbps UTP
1Gbps optical fibre
100Mbps UTP
SuperJANET connection via EastNET (~12.5 Gbps total)

Network

- Wireless:
  - Group has 2 wireless hubs serving it at present – our own "cheap and cheerful" Buffalo hub, and a departmental 3COM hub as part of an "official" rollout of wireless
  - Use limited to registered laptops and PDAs
  - 21 PC and 8 MAC laptops, 4 PDAs registered in hubs
  - Range is limited (~10-15m, which helps with security!) due to metal framework of building, so not all group members can see the hubs, and it is proving very difficult to find the best locations for them.

Video Conferencing

- Rather old (5 years) “mid-range” system (Zydacron Z360 (H.323) and ZC206 (ISDN) cards)
- Sony EVI D31 Camera (pan, tilt, zoom)
- hosted by (equally aging!) 500MHz Pentium III PC
- use (existing) data projector to display video on projection screen
- OK for up to ~12 people (though best for 6 or fewer!)
- We need to replace the facility soon, though it isn’t obvious that a room-based system is the best way to go nowadays.

- Hope that department would provide a VC system is as yet unfulfilled…

Software

- Nothing special…
  - AutoCAD, CADENCE etc for mechanical and electronic CAD work – all CAD work now PC-based .
  - Moving to SLC4 on desktops over the next few months.
  - Departmental license for Mathematica – allows home use for no extra cost.
  - Also run Maple, and MATLAB bulk deal is currently being negotiated.
  - XFree86 on Windows XP for X11 provision.
  - Group continues to run its own mail server (using Exim)

  - Condor used for batch as mentioned earlier.
  - Also have now subscribed to MSDN-AA to support the Windows GRID work.

Future plans and Concerns

- Cycle of desktop replacement is slowing – even 3-4 year-old PCs are adequate and the performance of new PCs is relatively static at present. Also increasingly problematic trying to run experimental software on desktop machines - so we're not under so much pressure from that source to buy new desktops.
- Extra disk space will (always) be needed: ~5TB/year is current best guess.
- Continue to enhance our farm provision – partly by taking advantage of CamGrid for HEP use.

- Continuing concern is how we manage all the extra kit in the medium term – especially as the system management team are heavily involved in ATLAS/LHCb commissioning.
- Security obviously a permanent concern. Department now has a firewall, but that doesn't help much with ssh attacks…