Royal Holloway
University of London
Centre for Particle Physics

UK HEP System Manager Meeting

30 June – 3 July 2009
RAL

Diego Bellini

Royal Holloway
University of London
Centre for Particle Physics

System Admin Staff:
(new admin staff member)

- Simon George
- Barry Green
- Diego Bellini (me)
- Govind Songara: Grid support (started around March)

Royal Holloway
University of London
Centre for Particle Physics

Some General Data:

- ~ 11 PP Academic Staff
- ~ 15 PP Research Staff
- ~ 11 PP Research Students
- ~ 20 desktops (SL4/SL5)
- ~ 21 servers (SL4, Win 2003, Solaris 8/10)
- ~ 5 laptops at disposition to all the Group
- ~ 32 personal laptop for research purpose

Grid Systems:
Faraday and Newton

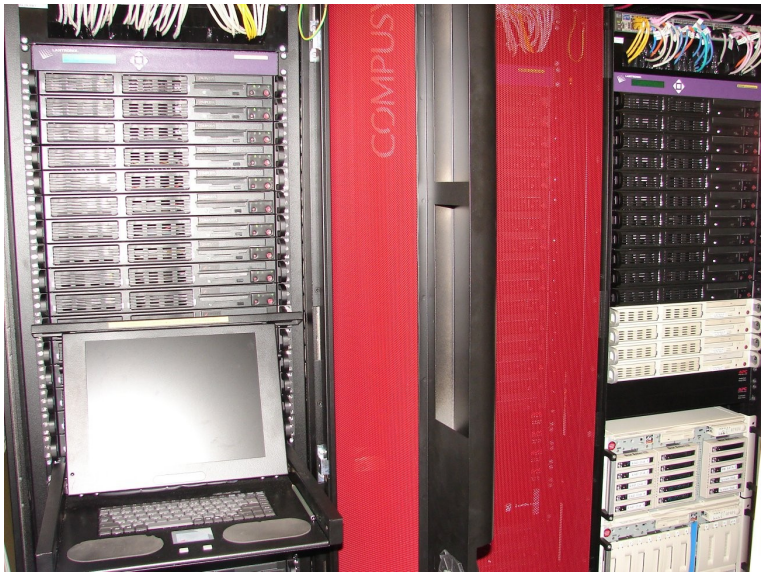

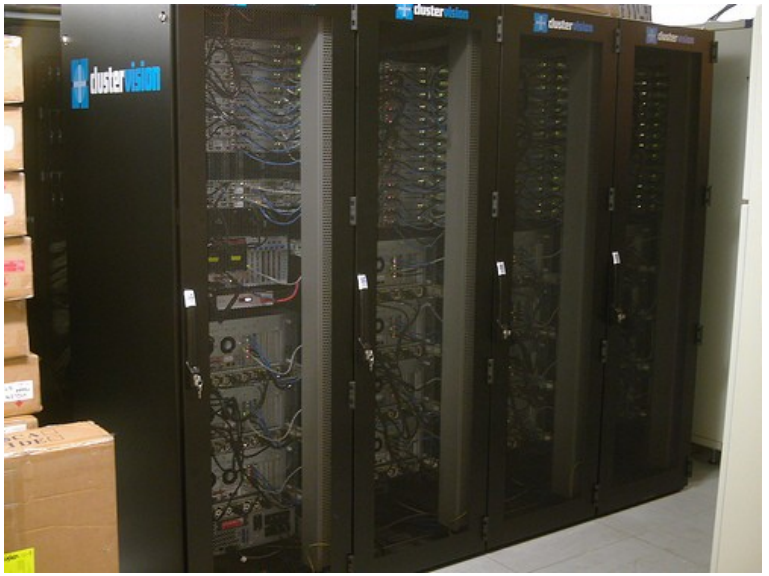

Grid Systems (cont'd):

A. Faraday

- Supplier: Compusys
- Location: RHUL PP computer room
- 5 years old – support ended March 2009
- 75 worker nodes with Intel Xeon (i686) 3GHz
  x2 cpu/node, 1GB mem/node
  => 163 kSI2k
- 3 storage nodes, total 8 TB RAID6.
- o/s: RHEL3
- Upgrade plan: SL4
  (since gLite i686-SL5 not supported)

Royal Holloway
University of London
Centre for Particle Physics

Grid Systems(cont'd):

B. Newton

- Supplier: Clustervision
- Installed December 07
- Current location: hosted by Imperial College
- Relocating to new RHUL machine room August 09
- 50 worker nodes with 2x Intel E5345 2.3GHz (Clovertown) (total 8 cores), 16GB mem/node => 960 kSI2k
- 15 storage nodes total 320 TB RAID6
- o/s: SL4

C. Next system: CIF/GridPP funded, planned for 2010.

Grid: Mini Test Cluster

- From 8 nodes of Faraday
  built up test cluster with SL4 32 bit
  for training & testing purposes

- Installing CreamCE, lcg-CE, bdii, torque,
  maui, worker nodes, etc. …

- Testing different monitoring and management tools
  like Ganglia, cfengine, nagios, MonAMI,
  etc....

- Will use to pilot SL4 upgrade of full cluster

ATLAS analysis (Tier3)

- Now: System for local interactive use and Grid UI.
- Future: share of next cluster

- Specified and ordered:
  - 8 core processing node with 16GB memory (Nehalem)
  - 60TB nfs scratch space for data

- New dual CPU system

- Vibration problems with 2TB disks

Power incident at IC

- Friday 16 May 2009 8am:
  Aircon failure in ICT machine room
  - Room overheating rapidly
  - Operators need to cut power
  - Go down aisles pulling power connections (IEC-C13)

- How we saw it
  - ssh prompt froze (Duncan and Simon were connected)
  - Nagios errors
  - RHUL network manager noticed firewall/router down
  - Rapid response from our IC contact explaing the situation.

Power incident at IC (cont'd)

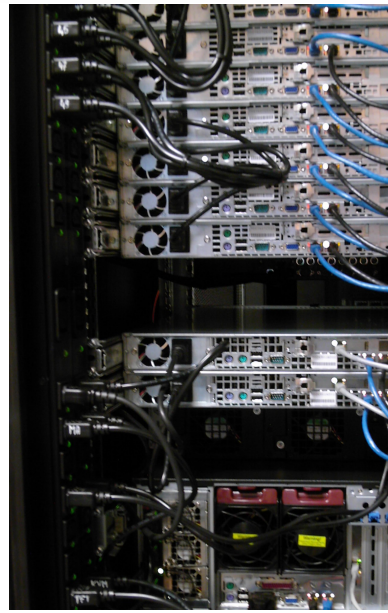

- Following Monday, Govind, Duncan and Simon discovered the full horror & spent a few hours reconnecting power cables.
- Thank goodness for:
  - Clustervision cable labeling & db on master node
  - Rack with broken lock: door could not be opened so remained cabled as an example.
  - Needless to say we are reviewing procedures.
  - CV/APC integration makes it easy to test a node is connected in the right place by powering on/off.

Since last September:

Ranking group web server
(www.pp.rhul.ac.uk)

- google search engine oriented

- google analytics tool

- distinguish web site areas

- page title, meta tag, site-maps,...

Windows server

- upgraded old win 2000 to win 2003
- authentication managed through Samba
- eliminated active directory

Started migration from SL4 to SL5

- desktop: planned to complete by the end of August

- server: rolling upgrade over next year, already underway

- cern-castor-client:

  group not recognized in kickstart's %packages section for i386 installation=> install in a later step through “yum groupinstall” (is there any motivation?)

Migration to outsourced spam filter (webroot)

- in progress

- some problems due to specific Particle Physics situation:

  - own DNS service for @pp.rhul.ac.uk

  - own mail service

File server migration

- /home nfs server
  - hardware: pc system, 2 Ethernet 1Gb (trunked), raid 5
  - o.s.: Solaris 10 with zfs (zraid)

- Backup system
  A. first file server
  - hardware: pc system, 2 Ethernet 1Gb (trunked), raid controller as jbod
  - o.s.: Solaris 10 with zfs (zraid2)

  B. second file server (off-site)
  located at new RHUL off-site machine room
  o/s & config to be decided
  - storage of dump file → no more tapes

File server migration

- Problems
  - slower performance if NVRAM is attached to raid controller
    We've tried to investigate but we didn't understand, anyway for our purpose it's good enough

  - switching to jbod raid controller feature:
    Raid controller couldn't reboot => upgraded firmware

  - reboot randomly doesn't work properly
    => open call with sun support

Plans for the future:

- Complete desktop upgrades
- Complete file server migration/new backup
- Complete spam filter migration
- Upgrade video room from fedora 8 to 11
- upgrade desktop hardware:
  - about 5 year old
  - dual core 64 bit: AMD Athlon 64 X2 5200+
  - ram 2Gb

Email migration

- We run a simple smtp & SSL imap server

- MTA: exim

- CC encourage us to move to their central Exchange2007 system

- Our academic staff want a shared calendar integration solution
  => motivation to move to them

- We are doing some trials now

Email migration – open issues

- migration will cost us several x annual effort running exim

- policy and technical issues with forwarding from CC to external address

- Anyone decent linux client for Exchange2007 over MAPI?

- Variety of personal calendar solutions already in use, complex syncing may be needed