

Imperial College London HEP Group Site Report

HEP Sysman June 2008

Kostas Georgiou

General Hardware

12 nodes for general access (batch, email, etc.)

- RHEL4 (33% i386, 66% x86_64)

~15TB storage for experiments/local users

Various servers

- NFS, Kerberos, LDAP, Web, Email, CUPS, Monitoring, SGE, etc.

2 Windows 2003 Domain Controllers

1 Windows 2000 Terminal Server

Grid Hardware

Clusters

- 50 2GHz Woodcrest nodes
 - » 4GB RAM, dual socket, dual core
 - » Access to local users but not to home directories
- 40 1GHz P3 nodes
 - » 1GB RAM, dual socket
 - » Reaching retirement age

~50TB storage in dCache

- 8 pools with external SCSI
- 4 older pools with 3ware cards

20 nodes for PPS, tests, etc.

- Mostly ancient hardware

Desktops

~100 Windows XP machines

- Still no plans for a Vista “upgrade”

?? Laptops running XP, Fedora, OS X, Vista

- Not managed centrally
- We have to monitor mac addresses in the network to find out when laptops are gone so we can reallocate ip addresses

~15 Linux desktop machines (i386, x86_64)

College is slowly taking over desktop purchasing

- Allocation is not always logical
- Hardware not always compatible with Linux

Infrastructure

Kickstart for all installs

- PHP script generates the kickstart on the fly
 - » http://...../ks.php?arch=x86_64&v=centos5&desktop&.....
 - » Needs to be rewritten

Puppet for configuration management

- Not every service has moved to puppet yet

Backups on disk

- rdiff-backup and/or rsync

Local rpm repository for software

- Everything is rebuild from source locally
 - » mock for the rpm builds

LCG

- tarball installations for WNs and UIs
- Lack of documentation always an issue

New Hardware

2 new servers with 1 6TB iSCSI box

- Xen virtual machines floating between the two servers with their “disks” on the iSCSI box
- Automatic failover by cman and rgmanager
- Home directories and mail server to be migrated there next week

5 new pools for dCache

- 24 x 1TB disks with Areca cards
- Arriving early next month

?? new worker nodes for LCG

- We will start looking for quotes in the next few days
- Most likely dual quad core boards with 2GB per core but you never know

Plans for the future

Finish upgrades to RHEL5

- Some lcg software will have to remain on older versions due to the glacial development speeds
- Start thinking about RHEL6

High availability

- Remove every single point of failure (or at least reduce them)
- Move all configuration to puppet
 - » The ability to rebuild any server/service from scratch in a matter of minutes is the key point here
 - » It can also serve as documentation

Investigate alternatives to XEN

- Dom0 support was dropped in Fedora 9
- It might come back in F10 or more likely KVM + xenner will replace it
- RHEL6 might not follow the Fedora route but it's very likely so it's best to be prepared