

High Energy Physics Group
Imperial College

UK HEP System Managers Meeting
1 July 2004

- 1 E450 Sun as a Mail/NIS/Web/Database server, 1TB disk space for experiment use and Home directories shared with NFS/samba
- 4 dual Xeon (2GHz) PCs for general use.
- 2 Windows Domain servers host the Active Directory and User Home directories (DFS)
- 1 Windows Terminal server for the Unix users and programs with limited licenses (e.g. Illustrator)
- 2 new storage servers to arrive next week (dual Opteron; 3TB each). 2 old servers will be used for backups (no more tapes)

- BaBar
 - 40 dual PIII (Coppermine 800MHz) nodes
 - 2 masters with 300MB space each
 - 2 storage servers with 1.4TB each
 - 1 Sun E450 with 1TB of space
- CMS (see Grid as well)
 - 40 dual PIII (Coppermine 1GHz) nodes
 - 5 masters with 400TB each

- Viking (shared resources at LeSC)
 - Currently ~520 CPUs (all >2.0GHz P4s), another 400 currently being added (Opterons)
 - 10TB disk space another 8TB being added
 - HEP share 33% (other groups computational engineering, bio-informatics)
 - Used for:
 - CMS Production (~15% of all CMS MC generated at Imperial)
 - D0 MC Production
 - LHCb production
 - Will be shared via LCG/EGEE soon (porting to SGE underway)

- EDG (finished March 31st)
 - Ran a Workload Management Service (RB) node on the production testbed
 - (Variable number of) CPUs provided for the production and other testbeds as required
- Currently
 - 66 CPU cluster on LCG2 (1.0GHz PIII)
 - 400GB Storage on LCG2

- August 2003: Our new computer room is ready
 - CMS/Grid farm moved to the new room. The Air-Conditioning unit can not handle the load. Farm is shut down
 - AC unit "fixed". We are told that some of the cold water pipes were connected the wrong way around
 - Several days later the AC-unit fails again. We decide to move the farm back until it's fixed
 - It's the cold water pipes again. Move to the new computer room is postponed until January
- January 2004
 - AC unit seems OK. All farms and servers are moved to the new room

- February - March 2004
 - Tuesday 26th late afternoon
 - Temperature goes over the limit and the power to the whole room is cut. Nobody available at the time knows what to do
 - Wednesday 27th - Thursday 4th
 - The Engineers replace the drive belts in the AC and restore the power. Temperature is high and keeps rising. We try to alert people with no success.
 - Friday 5th late afternoon - Monday 8th
 - AC fails again. Since we can't find anyone that knows what to do we shutdown everything leaving only essential services running and we go home for the weekend
 - Monday 8th
 - AC is repaired.

- About 10 Linux Desktops running RHEL3 so far
 - Many useful applications are missing e.g. CD burning, mp3 players etc.
 - Software updates through College RHN proxy server
 - We have our own channels at the proxy to provide/upgrade HEP specific rpms
- RHEL4 to arrive soon (Beta1 is already out).
 - Do we want to upgrade again?
- Fedora 2 has many advantages over RHEL
 - 2.6 kernel, SELinux etc..
 - The fast (6-8 months) upgrade cycle is a good thing for the Desktop. People want the latest browser version etc.

- Around 80 desktops and 30 laptops
 - The time that we'll have more laptops than desktops isn't too far away
- No new Windows 2000 installations, Windows XP is now the default
 - Longhorn won't arrive for some time (~2007) so no more upgrades in the near future
- College is now providing new desktops to staff members
 - Part of the College domain and not our own
 - Administration partly from College staff
 - Registered owner of the computer has Admin rights

- Wired
 - No more hubs at our network
 - We need to start buying more gigabit switches soon.
 - All new servers use gigabit network cards (some more than one). A few desktops/laptops have gigabit as well.
 - iSCSI seems promising (more gigabit ports)
- Wireless
 - College wants to take over all APs
 - You need a College account and a PPTP client to be able to see the outside world
 - We run our own system based on NoCatAuth
 - Only requires a web browser to work

- Security still a major problem
 - Farms need outdated versions of RedHat; same for the Grid. Something has to be done
 - Hard to ensure that laptops keep up with security patches
 - SPAM is increasing

Future plans

- NIS has to go
 - Kerberos/LDAP is the obvious solution
 - Integration with Active Directory?
- NFS needs to go: insecure, doesn't scale, SPOF
 - NFS4 not support in all systems
 - Cluster Filesystems
 - RedHat GFS is open source now
 - High Performance
 - High Availability
 - Linux only
 - TerraGrid from Terrascale
 - Costs £££
 - Linux only
 - Others?

- P4 is often slow compared to P3/Athlon for codes that aren't optimised specifically for it
 - For example, the following code (compiled with gcc 3.3.3):

```
#define ISIZE 200
#define JSIZE 200
double x[ISIZE], y[JSIZE], a[ISIZE][JSIZE];

for( i=0; i<ISIZE; i++) {
    for( j=0; j<JSIZE; j++) {
        x[i] = x[i] + y[j] * a[i][j];
    }
}
```

- Average time to complete in seconds (lower is better)

