

Winnie Lacesso

Bristol Site Report

June 2010



Liveware

- Departmental / Administrative / Teaching / Lecture Theatre / Lab Computing / Networks:
 - Dr JP Melot (Networks & Microsoft guru)
 - Neil Laws (Microsoft / hardware guru)
 - Dr Rhys Morris (Astrophysics support)
 - Winnie Laccesso (.5 PP support)
 - Dr Bob Cregan (.3 PP support): mainly IS HPC Storage support
- About 40 PP staff & students
- 286 staff in School of Physics (*was: Department of Physics*)
- PP Tech support works closely with Bristol IS HPC Admins

PP Hardware & UBris IS HPC

- About 80 desktops/laptops (many staff have 2 to 4): about ½ are Macbooks , few SL (ScientificLinux), rest XP
- 19 SL4/5 non-WN servers: NFS, AFS, compute, swe.lic, etc
- VMWare Server 2.0 on 2 x SL5 svrs, 4 production VMs (SL4/5)
- 1 Windows Server 2003 (fileserver)
- Bristol HPC (Managed by IS): #379 in 2010 Top500! (IS HPC Admins surprised Bristol still on the list!)
- HPC Stage1 = SL4 400 cores; HPC Stage2 = SL5 4000 core
 - Both get very heavy usage
 - Both use gpfs as shared filesystem (home dirs & 11TB cluster data storage) over WN;
 - Both see patchy very poor gpfs performance (user complaints!)
 - Large data storage = few hundred TB = gpfs multiclustered with HPC gpfs cluster (PP has 109TB managed by LCG SE)


PP Servers: Functions

- Production SL 4 AFS server (/afs/phy.bris.ac.uk) replaced Win2000 antique AFS server. AFS = odd-ish beast, internal backend DB, to user it look like Kerberos-authenticated NFS. Works well except when it breaks..... then is byzantine weirdness
- Win2003: fileserving (480GB); considering Unix/Samba replacement
- Rest of PP Linux servers = SL4/5:
 - 2 NFS with 11TB (one also RO Samba)
 - PBS batch(3) (*retiring, to be replaced with Departmental*)
 - Compute (~3)
 - web/subversion/elog, mediawiki
 - infrastructure (DHCP, kickstart)
- PP gets 90 jobslots on SL4 HPC cluster (2GB RAM/core)
- Not yet using SL5 HPC cluster (only 1GB RAM/core)

Storage Story

- CMS fast-I/O server: DNUK Teraserve R2450x 2U, 2 x Intel XEON E5405 (2.0GHz), 16GB RAM, 6 x 1.5TB disks. Bought July 2009.
- 50GB / on /dev/sda, 2GB swap on each sda..f, rest of sda..f = RAID0 /storage
- Gives 7.9TB usable /storage
- Remind users repeatedly: NO backup, NO redundancy; disk error = **GONZO**. *Nothing in /storage should be unique – must be copy of something elsewhere* (this was mangement's decision)
- 26 May 2010 disk failure – bits of RAID0 /storage start to vanish!
- Notify users; have spare 1.5TB disk – have to destroy & recreate RAID0 /storage
- *Some users still panic'd & frantically tried to copy /storage data to NFS* (They didn't 'fess if it was unique... maybe just troublesome to recreate)
- Moral of story: Remind users frequently; Lie &/or Invent errors to provoke users to behave in safe way??

Gridftp01 under pressure(not...)



- maillog is a good canary:
- May 3 00:05:17 gridftp01 sendmail: rejecting connections on daemon
MTA: load average: 2763
- `root@gridftp01> ps -ef | grep gridftp | wc -l`
3275

Security

- User laptops frequently go offsite (home, CERN, RAL), come back & reconnect to internal network. No (detected) incidents. Even from users with root/admin access on laptops.
- One laptop stolen on train while student reaching up to retrieve his pack from rack above seating.
- Fortunately, USB backup disk kept in *different* location. Moral of story: carry USB backup disk separate from laptop.
- University considering requiring Full-Disk Encryption for laptops or removeable media with sensitive data
- All PP staff have said no sensitive data on laptops / USB
- All XP SP2 will be upgraded to SP3 this month

Issues

- **Upcoming/pending work :**
- VMware has good & bad points; is EndOfLife June 2011, so need to learn about + migrate to KVM (RHEL is ditching Xen)
- Retire old hardware – anyone want some 32-bit WN?? Or PATA 3ware RAID cards? *Can't give it away!*
- **Recent/ongoing problems :**
- UPS needs rearranging (*again!*) – some important servers not on UPS
- Funding 100% doubtful
- Senior Network / Compute Manage JP Melot retiring July 2011 at latest (possibly sooner) – no replacement.... =:o