

Local Computing Issues

Group Meeting, 8th September, 2014
Mark Slater + Matt Williams, Birmingham University

Since Lawrie left things have *mostly* been running OK with a couple of glitches here and there as I figured out the systems

I have also tried to respond quickly to problems and requests but this should improve after I've finished the move of all the services

The current state of the local systems:

- Batch Cluster Farm, 192 cores, 2GB per core ●
- 120TB Storage ●
- Fedora 18/20 Desktops ●
- SL5 and SL6 images ●
- Two F20 login nodes ●

Please let me know if anything is missing/not working or there's some software or service that you want!

General service changes in the last 12 months include:

- ~3/4 of the systems have been transferred to Puppet ●
- Nearly all desktops upgraded to Fedora 20 ●
- Added CVMFS service ●
- Transferred CUPS and Torque to Virtual Machines ●
- Moved the SL5/6 images to live Virtual Machines ●

This leaves the following services still to be updated:

- The mail system ●
- Web pages and server, both user and group ●
- Windows Terminal Server ●

Worker Nodes

From Kostas' EU grant, we purchased 2 new worker node boxes that replaced the aging and flaky nodes. These provide 192 cores but I plan to add more memory and HD space to allow 256 cores.

Service Nodes

To allow services to run on dedicated Virtual Machines, I've bought 2 new servers that will be able to host current and future VMs

Mass Storage

Currently, almost all the storage we have is out of warranty so I have ordered 160TB of new storage which has arrived but needed to be taken away again (no HDs!). Hopefully this will be added in the next 2-3 weeks.

Home Area

The home area is also out of warranty and a replacement is currently being set up. I plan to do the migration at the end of September.

Windows Services

The current Windows Terminal Server is in serious need of replacement. There are 3 options:

- Replace with the latest version of Windows Server (~£500)*
- Provide an option to login to University Citrix server (~free)*
- Do both*

The Citrix server provides a desktop with Office, University drives and a few other programs so would be as good as the WTS. However, we couldn't add our own stuff to it and if too many people use it they may come asking for money!

Security, Logins, etc.

At present, we only allow logins to eprexa/b from select IPs via username/password. To better comply with future uni security regulations and to reduce my blood pressure, I would like to introduce a two factor method. Either:

- SSH key login only allowing those keys generated by our system*
- Certificate based logins*
- Google App based 'proper' two factor login*

I would propose the first as this is the least work for everyone. The other two methods are more secure but require more 'faff' on the user end.

Batch System Upgrade

We currently use the Torque/Maui system for running jobs on the cluster. This is (very) old, out of support and causing problems. I plan to switch to a different system (Condor) that is a lot easier to run. There will be changes required to submission scripts but I'll try to minimise this as much as possible

Changes to Mass Data Storage

At present, the RAID disks are NFS mounted and individually named. This is a bit awkward both for users and for quotas, etc. When the new storage has been returned, I'm going to install it using a distributed system called Lustre. This will present all storage as one mount point (/lustre).

Server Room Clear Out

The server room is in a bit of a mess and with all the new equipment, the air con is not running very efficiently. I plan to have a big clear out and condense all the servers down to 5 racks that will be placed along the left wall. This should improve air circulation.

Since Last year, I've taken over most of the management of the Grid site.

We've been looking at improved monitoring and replacing some of the older systems to make things more maintainable.

We've added:

- 12 new 32-job compute nodes. Now at 1200 job slots ●
- ~300 TB of storage taking us to ~650 TB ●
- SL6 on all systems ●

I'm investigating a system to make use of unused cluster time to process extra grid jobs.