

QMUL Site Report

Cozmin Timis

Particle Physics Research Centre

The Team

- Alex Owen – Snr. Physicist/Programmer
- Cozmin Timis – Physicist/Programmer
- Dan Traynor – Snr. Grid Site Admin
- Terry Froy –Grid Site Admin

The Team

- Major changes in IT team structure at school level.
- Alex promoted in managerial role, the two PPRC Physicist/Programmers to take more school duties.
- At School level: 2x Physicist/programmers, 1 web technologist, 1 Sysadmin (leaving QM in 2 days. Please encourage people to look at advert on jobs.ac.uk or <https://webapps2.is.qmul.ac.uk/jobs/job.action?jobID=3418>)
- Double the users numbers, different modus operandi
- Huge number of servers/services to deal with on disparate hardware
- Large numbers of windows desktops to look after
- Good thing: ProxMox cluster setup

Physics Server Room

- 8 racks (new one added for school computers)
- 32xA rack (2x16A) with 2 UPS per rack
- Max Input Power= $32 \times 8 \times 220 = 56 \text{kVA}$
- 4x14KW Aircon => Max Cooling 56kW
- 20Gb/s (2x10Gb) to server room with 2GB/s (2x 1GB/s) backup link
- APC Temp sensors on UPS card



PPRC Computing

- No new server or WN hardware
- Extended Dell server warranty from resellers (Esteem)
- About to purchase ~ 130TB of storage

PPRC Computing

- 8 Academics
- 16 "Postdocs"
- 18 PhD Students
- Total 43 People

	Lin	Win	Mac	Tot
PC	39	15	0	54
Lap	11		14	25

Servers	6: LTO5 backup
Batch	544 Job Slots
Disk	233 TB

Astro Computing

- 18 Academics
- 8 Postdocs
- 18 PhD Students
- Total 43 People

	Lin	Win	Mac	Tot
PC	40	2	6	5
Lap	13		14	27

Servers	5: LTO5 backup
Batch	96 Job Slots
Disk	102 TB

Additional kit

- 3 VR Alienware laptops + Oculus Rift sets
- DAQ Test Rig. Pair of Dell R520 Netgear XS708E (8 port 10GBaseT)
- CUDA: K20m(Astro) GeForce GTX 1080 (Tensor Flow) Tesla P4
Passive to come

OS Versions

- Debian Servers: Mostly Debian 8
- SL servers SL6
- Desktops: Debian 8 / SL 7
- SL6 batch queues
- SL7 now the default option

Laptops / PC from ITS

	Laptop	PC
Academic	PPRC/ITS	PPRC/ITS
Postdoc	PPRC/ITS	PPRC/ITS
Postgrad	Possible ?	PPRC

Possible SL7 desktop solution from ITS ?

GridPP at QMUL

- 2*FTE 100% GridPP focused (Daniel Traynor + Terry Froy) + ~0.1 FTE unpaid support from others.
- ~5PB storage (Lustre 2.8).
- ~4000 job slots (1/2 with 2GB RAM/slot, 1/2 with 4GB RAM/slot). 6 GPUs (K40 class)
- 10Gb/s SFP+ to every node, multiple 40Gb/s from core to top of rack switches. 20Gb/s WAN link.

GridPP at QMUL



looks something like this

New Storage for GridPP cluster

HGST 4U60G2 JBOD

60*10TB SAS disks

2 HPE DL380, 2*40Gb network

Dual connect JBODs to servers
provides Lustre failover

1/2 rack, 2400TB RAW, 1.6 PB,
24 10 disk raidZ2 arrays



New Compute for GridPP cluster



HPE DL385,
AMD EPYC 7351 (2.4GHz)
64HT cores, 256 GB RAM
HepSpec 740

Too new for good SL6 support:
legacy BIOS PCIE PXE boot broken
and kernel core topology wrong



Dell R440,
Intel Gold 5118 (2.3GHz)
48 HT cores, 192 GB RAM
HepSpec 575

New Physical Infrastructure for GridPP

Replacing old generic racks
with APC AR3300



Replacing failing PDUs
with eaton ESWB16



New Software Infrastructure for GridPP

- Three node Proxmox VMhost cluster with failover using ceph storage.
- Several bespoke CentOS 7 services set up (monitoring, lustre storage server, Dev and test servers).
- New Deployment system with failover: custom database + Cobbler (kickstart), Ansible (configuration), PowerDNS and ISC Kea DHCP services. To be used for bulk CentOS 7 deployment.
- Development Hadoop and Spark mini cluster. Work ongoing to integrate Hadoop/Spark with main cluster e.g. using magpie (<https://github.com/LLNL/magpie>).

Plans for GridPP at QMUL

- Continue to replace all racks and PDUs.
- Preparing for 100Gb/s WAN link for run 3. Storm /gridftp/webdav servers with 100Gb/s capability.
- Long term: Investigate Openstack/Kubernetes/ceph platform to run under/instead of the grid software.
- Continue to push for refresh of data centre, prepare for possible move to new one.
- Update from NIS to LDAP/Kerberos.