



GridPP

UK Computing for Particle Physics

RAL Site Report

HEPiX Fall 2013, Ann Arbor, MI

28 Oct - 1 Nov

Martin Bly, STFC-RAL



Science & Technology Facilities Council

e-Science



- CPU: ~97k HS06 (~10k cores)
- Storage: ~8PB disk
- Tape: 10k slot SL8500
- FY13/14 procurement
 - ~7.0PB useable capacity disk storage
 - ~46k HS06 CPU
 - Tenders evaluated, in EU standstill period
 - Much better benchmarking by vendors this time!
 - Same as the FY12/13 model with bigger drives or faster CPUs
- 2007 kit mostly decommissioned
- 2008 generation being phased out

- WAN
 - RAL site has migrated to new Janet6 (aka SuperJanet 6) backbone
 - Dual 30Gb/s active/passive failover link
 - Two routes on to site
- Tier1 link to boundary now re-established at 20Gb/s
 - 10Gb/s to CERN, 10Gb/s to Janet6
- LAN
 - Mix of Dell/Force10 S4810p & 60, Arista 7124, Nortel/Avaya 55xx & 56xx series. Mesh and routing with z9000 and Extreme x670
 - Tier1 migration to Mesh network
 - Testing internally complete, testing routing connections

- HTCondor selected as replacement for Torque/Maui
 - Rollout in progress
- Migrating from CREAM to ARC CEs
 - SL6
- Current status
 - All batch resources on SL6
 - 50% CPU resources moved to HTCondor
 - Migration will be completed in early November
- Testing opportunistic use of resources from private cloud in the batch system
- See talk by Andrew Lahiff

- SL6 migration mostly done
- New FTS3 system
 - MySQL backend database, VMs
 - Advanced testing
 - Production transfers with Atlas
- Quattor/Aquilon
 - 200 systems now managed with Aquilon
 - Research Infrastructure Group experimenting with it
- Most services on VMs
 - Have or had issues with ganglia, BDIIs
- CVMFS service working well
 - Talk on Stratum-0 by Ian Collier

- Stats as of October 2013:
 - 64m files
 - Stored data capacities: 14PB on tape and 8PB on disk
- Recent news:
 - Preparing for T10KD tape media for production next year
 - Developed a new WebDAV front end
 - Not fully production ready yet.
- Next generation (disk) storage project continues
 - Challenge: Try to find something simple and easy to run in conjunction with CASTOR for tape access
 - Review of options in May concluded that there was no compelling reason to move from CASTOR at this time
 - Testing of CEPH continues as option for storage in Cloud infrastructure
- AFS: Terminating RAL cell on November 5th

- Uprating of ‘Essential Power Board capacity’
 - 400A to 630A
 - Requires complete isolation from supply and UPS
 - No UPS supply to protect HA services etc
 - Tier1, HPC and Corporate systems
 - Plan to have minimum service running
 - Migrate various essential Corporate services up to Daresbury Lab
 - Migrate Tier1 VMs to Hypervisor cluster in old computer centre
 - Shutdown Castor, batch, all non-essential services
- Mandatory testing requirement - regulatory requirement
 - Taking opportunity to test all UPS circuits from source
- Half-day for Essential Board, Castor fast-restart thereafter
- Batch resumes when all other services are back and stable

- ‘Facilities’ are using Castor, StorageD services
 - Run by SCD
 - Dedicated Castor instance with a set of associated services
 - Separate Nagios instance
 - using "Icinga" rather than "Nagios"
- Various UPS generator ‘failures’
 - Failed starts, failure to assume load etc
 - Due to latent faults and poor initial installation
 - More rigorous testing regime
 - Full load tests conducted more often

