



GridPP

UK Computing for Particle Physics

RAL Site Report

HEPiX Spring 2014 Annecy-le Vieux

19-23 May

Martin Bly, STFC-RAL



Science & Technology Facilities Council

e-Science



- CPU: ~127k HS06 (~13k cores)
- Storage: ~13PB disk
- Tape: 10k slot SL8500 (one of two in system)
- FY13/14 procurement
 - CPU: 32 x Supermicro Twin², 2 x E5-2650v2, 128GB RAM, 2 x 2TB HDD
 - Storage: 57 x 36-bay Supermicro chassis, ~120TB useable per system
 - 34 x 4TB WD SE HDD / LSI 9261-8i
 - 36 x 4TB WD RE HDD / LSI 9271-4i
- 2008 generations being phased out
- 2009 generations phase out started
- FY14/15 procurement
 - Similar to last year, starting soon.

- Tier1 LAN
 - Mesh network enabled
 - Two Dell Force10 Z9000 in active-active VLT pair
 - 2 or 4 x 40Gb/s LACP to/from each S4810P
 - Some S4810P VLT pairs for resilience
 - Services transferring to it
 - Phase 1 of new Tier1 connectivity enabled
 - Routing to RAL Site now via active/passive pair of Extreme x670V switches
 - 20Gb/s redundant link from each
 - Phase 2: move the firewall bypass and OPN links to new router
 - Will provide 40Gb/s pipe to border
 - Phase 3: 40Gb/s redundant link to RAL Site
- RAL LAN
 - Migration to new firewalls almost complete
 - Migration to new core switching infrastructure almost complete
 - IPv6 test network soon
- Site WAN
 - Dual 30Gb/s active/passive failover link to Janet6

- **Batch system**
 - Migration from Torque/Maui to HTCondor completed in November 2013
 - Currently running HTCondor 8.0.6
 - Very stable operation, no major problems
 - Multicore jobs running successfully since November
- **CEs**
 - ATLAS & CMS only using ARC CEs
 - Gradually moving remaining VOs from CREAM to ARC
 - Aim to phase out CREAM CEs
- **Talk: Ian Collier ‘A year of Condor at RAL Tier 1’**

- SL6 migration (still) mostly done
 - Most services on VMs
- FTS3
 - A primary test site, extensive testing
 - Now a production instance
- Quattor/Aquilon
 - Talk: 'Quattor Update' - Ian Collier

- Deployment at RAL supported by GridPP
 - EGI Infrastructure
 - Initially for UK VOs
 - extended to international small VOs and 2 NGIs
 - Web interface for SGM to upload and unpack tarballs, and publish
 - New GSI interface to transfer and process tarballs
- 11 repositories published at RAL
 - Separate Stratum-1 service for non-LHC VOs
 - 160GB published on Stratum-0 @ RAL
- EGI CVMFS task force
 - KO meeting August 2013, regular meetings
 - promotes the use of CVMFS technology by user communities
 - Network of sites providing Stratum-0, Stratum-1 and squids

- Two production clusters with shared storage, several local storage hypervisors
 - Windows Server 2008 + Hyper-V
- Issues with VMs
 - Stability and migration problems
- Re-build the shared-storage clusters from scratch
 - New configuration of networking and hardware
 - Windows Server 2012 and Hyper-V
 - Currently migrating most VMs to local storage systems
- Aim to have three ‘new’ clusters
 - Include additional hardware with more RAM
- Talk: ‘RAL Tier 1 Cloud & Virtualisation’ - Ian Collier

- **Castor**
 - June: Upgrade to new major version (2.1.14) with various improvements (disk rebalancing, xroot internal protocol)
 - New logging system with ElasticSearch
- **Ceph**
 - evaluations continue
 - Talk: ‘Ceph at the UK Tier 1’ - George Ryall
- **Storage woes**
 - 1 of 2010 sets (18 x 36TB) to be decommissioned early
 - 4 catastrophic failures including 2 data loss over 2 years (20%)
 - 4,3,7,3 drives thrown nearly simultaneously
 - SM chassis, Adaptec 5405, WD 2TB RE4 WD2003FYYS
 - So far unable to isolate - backplanes or disks

- UPS ‘shutdown’ for circuit testing
 - Successfully completed in November
- UPS generator load tests
 - No further failures, test schedule reverted
- Windows XP ‘banned’ from site networks
 - Almost all desktops and laptops upgraded to Windows 7
- New telephone system rollout imminent
- Recruiting a grid-admin soon
- AFS: RAL cell was terminated on November 5th 2013

