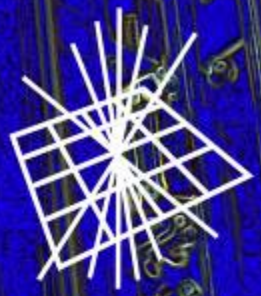

GridPP

UK Computing for Particle Physics

RAL Site Report

Martin Bly

HEPiX Spring 2009, Umeå, Sweden

Science & Technology Facilities Council

Overview

- New building
- Tier1 move
- Hardware
- Services
- Castor
- Storage
- Networking
- Developments

New Building I

- New computer building now ready
  - Originally due 1 Sept 08
    - Several interim due dates
  - Accepted Feb 09
  - Cleared for use March 09
- Delay mainly caused by building over runs
  - Generally building projects always take longer than scheduled
    - Unless you have lots of money
- Planning blight
  - When can we install the new procurements?
  - When can we move the Tier1?
    - Originally: during LHC shutdown in Jan-Mar 09

New Building II

- Oct 08: the LHC broke …
  - …which threw all the scheduling up in the air
  - Forge ahead with Castor upgrades, continue with procurements, possible handover of building in December, plan move for as late as possible in original LHC downtime schedule
- New LCH schedule announced
  - long run through 2009/10, starting early Autumn 09
  - Move still scheduled for Feb 09
- Move delayed indefinitely in mid-Jan when it became clear no firm date could be predicted for acceptance of the building
- Decide to delay Tier1 move to last possible window before experiments require computing services stable prior to announced LHC data taking schedule, with enough time to attain stability of Tier1 after the move
- Meanwhile, concentrate on upgrades necessary for long run because there may not be opportunities to do it after data taking starts
  - Not much time…

Tier1 move

- RAL Tier1 will move 22 June – 6 July
  - Subject to final go-ahead
    - Requires agreement of all team leads and service manager, machine room operations group (not Tier1 staff)
  - Building throwing up new problems every day
    - Most recent due to errors in power infrastructure and distribution systems
- Sequence is complicated by the need to reduce the data transfer services down time to a minimum
  - RAL hosts several services required for UK-wide data access systems (FTS, LFC, PROXY, …)
  - Regarded as unacceptable to have these down for extended periods
- Order is Batch, Services nodes and databases, storage, silos
  - Published in GOCDB (which is being replicated off site!)

New Hardware

- ~3000kSi2K (~1850 cores) in Supermicro 'twin' systems
  - E5420 / San Clemente, L5420 / Seaburg
  - 2GB/core, 500GB HDD
- ~2PB (110 servers) in 4U 24-bay chassis – 2 vendors, mix of:
  - single Areca and dual 3ware/AMCC controllers
  - Seagate and WD drives, 22 data in RAID6, 2 system in RAID1
- Second SL8500 silo, 10K slots, 10PB
  - In new machine room – pass-through to existing robot when relocated
  - Tier1 use – GridPP tape drives will be transferred
- Services nodes etc
  - 10 'twins' (20 systems), twin disks
  - 3 Dell PE 2950 III servers and 2 array units for Oracle RACs
  - Extra SAN hardware for resilience

Old Hardware

- Decommissioning program:
  - 199x:
    - Old AFS server replaced
  - 2002:
    - 156 x dual PIII 1.4GHz CPUs
    - Finally! - 44TB SCSI array storage out of use
  - 2003:
    - 80 x dual P4 Xeon 2.66GHz CPUs
      - Now used for testing
    - 40TB SCSI array storage
  - 2004 – hardware to be decommissioned soon:
    - 256 x dual P4 2.8GHz CPUs
    - 20 x 2 x 3.5TB (~140TB) SCSI array storage
- None being moved to new building

Service changes

- Original AIX/Transarc AFS server replaced by three new Linux/OpenAFS servers
  - RAID10, 1.4TB each
  - Project to use AFS to provide gLite software for WNs
- VO software areas
  - LHC VO moved to individual servers
    - CMS now 64bit OS due to size of RPM database
- Additional CEs to increase resilience
  - Service for individual VOs spread over more than one system
- RBs decommissioned, WMS/LBs commissioned (3+2)
- Additional hosts for LFC, FTS front end services
  - Oracle RAC for FTSLFC backend
    - 64bit, mirrored ASM
    - Additional faster arrays to be added for resilience

Batch Services

- Main service:
  - SL4.7/32bit
  - ~2200 cores from 2.8GH Xeons to 5440 Harpertowns
  - Torque/Maui on SL3/32bit server
- Next Service:
  - SL5(.3)/64bit with SL4 and 32 bit compatibility libraries
  - Torque/Maui on SL5/64bit server
  - In test for VOs
  - Rollout and resources as required by VOs – timetable not fixed
    - Would like to roll out before data taking, particularly for new hardware
  - Swing all 64bit cores to SL5
  - Probably retire 32-only nodes – no firm decision or commitments

Castor

- Problematic upgrade to 2.1.7 in late summer 08
  - Bugs, stability issues, long running issues only now being understood
  - Huge drain on Castor and Databases teams at RAL
    - As well as folk at CERN offering support
  - Much more stable since Xmas
  - As a result, very careful consideration given to subsequent upgrades
  - Enhanced testing at RAL
  - More conservative downtimes
- 2.1.8 considered for deployment before data taking
  - Extensive consultation with VOs: no overriding need for data taking
  - Time to do it and recover stability before Tier1 move, or do it after move and before data taking?
  - No, large scale deployment of 2.1.8 at RAL postponed
  - Testing will continue, to be ready if necessary

Storage

- ECC issues
  - Bios configurations
- Enhanced HDD failure rates
  - ~6%. Looks as if we may be running too cool!
- Firmware issues
  - Failure to alarm on certain modes of disk failures
  - Failure to start rebuilds
- Firmware on 3ware cards 'tuned' to ext3 file systems?
- Aim to deploy new hardware with SL5/64bit/XFS

Networking

- Force10 C300 in use as core switch since Autumn 08
  - Up to 64 x 10GbE at wire speed
- Implementing routing on C300
  - Easier routing to LHCOPN and bypass to site firewall for T2 data traffic
- Nortel 55xx series stacks at edge
  - CPS farm and storage attached to each stack
  - Experimenting with bigger stacks and trunked 10GbE uplinks to C300
    - Is combined uplink traffic less than twice the traffic for two single uplinks due to greater probability of data being on a server on same unit or stack?
- Relocation of Tier1 to new building will split Tier1 network until remaining legacy kit is retired

Developments

- Fabric management
  - Current system: PXE/Kickstart scripts hand crafted
  - Worked well for several years but now showing the strain of complex deployments
  - Castor team have been using Puppet for some castor-related extras
    - Gridmap files, config files
  - Time for a comprehensive review for potential replacements
  - Quattor top of (short) candidate list
    - Does the necessary
    - Lots of support in-community
    - Recipes for grid-type deployments available
  - Trial Quattor system using virtual systems
  - Plan to have a useable production deployment by late August
    - Deploy new procurements
    - Spread to services nodes and existing systems over time

Shameless plug

- See talk on the RAL New Building experience in the Data Centres session on Thursday afternoon:

http://indico.cern.ch/contributionDisplay.py?contribId=38&sessionId=21&confId=45282