


RAL Site Report

Spring HEPIX @ CERN
5-9 May 2008

Martin Bly



Science & Technology Facilities Council
e-Science



GridPP

UK Computing for Particle Physics



Overview

- New Building
- Site issues
- Tier2
- Tier1





New Computing Building

- Previous HEPiX: report on new computing building and how we went about procuring it.
- November 2007 - looked like a sparse Meccano construction - just girders
- Now has walls, a roof, windows, skylight - looks like a building
- External 'beautification' starting
- Internal fitting of machine room level yet to start
- On track - completion due autumn 2008
- Migration planning starting
 - Target: To move most of Tier1 hardware Jan-Mar 2009



5 - 9 May 2008

RAL Site Report - HEPiX @ CERN



Portable Device Encryption

- Big concern in the UK over data-loss by 'government', like everywhere else.
 - Mostly careless custodianship rather than 'enemy action'
 - Many stolen/lost laptops, CDs/DVDs going missing in transit...
- Government has mandated that all public service organisations must ensure all portable devices taken off their site have the data storage encrypted by an approved tool
- This means: all laptops and other portable devices (PDAs, phones) which have access to 'data' on the RAL network must have encryption before they leave site
 - 'Data' means any data that can identify or be associated with any individual - thus Outlook caches, email lists, synchronised file caches of 'corporate' data of any sort
- Many staff have rationalised what they keep on their laptop/PDA
 - Why do you need it? If you don't need it, don't keep it!
- Using Pointsec from Check Point Software Technologies Ltd
 - Will do Windows XP, some versions of Linux
 - ...but not Macs, or dual-boot Windows/Linux systems (yet!)
- Painful but necessary
 - Don't put the data at risk...

RAL PPD Tier2

- Part of the UK SouthGrid T2:
 - Birmingham, Bristol, Cambridge, JET/Culham, Oxford, RAL, Warwick
- Mainly supports LHC VOs, but also FNAL, SLAC, DESY, EGEE etc
- Also support for all local particle physics groups
- Takes advantage of RAL Tier1 procurement tendering
- CPU: ~920kSI2K (1/3 is recent)
- Disk: 160TB

Tier1: Grid Only

- Non-Grid access to Tier-1 has been ended. Only special cases now have access to:
 - UIs
 - Direct job submission
- Until end of May 2008:
 - IDs will be maintained (disabled)
 - Home directories will be maintained online
 - Mail forwarding will be maintained.
- After end of May 2008
 - IDs will be deleted
 - Home filesystem will be backed up
 - Mail spool will be backed up
 - Mail forwarding will stop
- AFS service continues for Babar (and just in case)



CASTOR / dCache


- CASTOR: production version is v2.1.6-12 hot-fix 2
 - Recently much more stable and reliable
 - Good support from developers at CERN - working well
 - Some problems appear at RAL that don't show in testing at CERN because we use features not exercised at CERN - speedy investigation and fixing
- Considerable effort with CMS on tuning disk server and tape migration performance
 - Recent work with developers on migration strategies has improved performance considerably
- dCache service closure announced: end May 2008
 - Migration of data proceeding
 - Some work to do to provide generic Castor instance for small VOs
 - Likely the closure deadline will extend some months

Hardware: New Capacity Storage

- 182 x 9TB 16-bay 3U servers: 1638TB data capacity
- Two Lots based on same Supermicro chassis with different disk OEM (WD, Seagate) and CPU (AMD, Intel)
- Dual RAID controllers - data and system disks separate
 - 3Ware 9650SX-16ML, 14 x 750GB data drives
 - 3Ware 9650SX-4, 2 x 250GB or 400GB system drives
- Twin CPUs (quad-core Intel, dual-core AMD), 8GB RAM, dual 1GB NIC

Backplane Failures (Supermicro)

- 3 servers “burn out” backplane
- 2 of which set off VESDA
- 1 called out fire-brigade
- Safety risk assessment:
Urgent rectification needed



- Good response from supplier/manufacturer
- PCB fault in “bad batch”
- Replacement complete

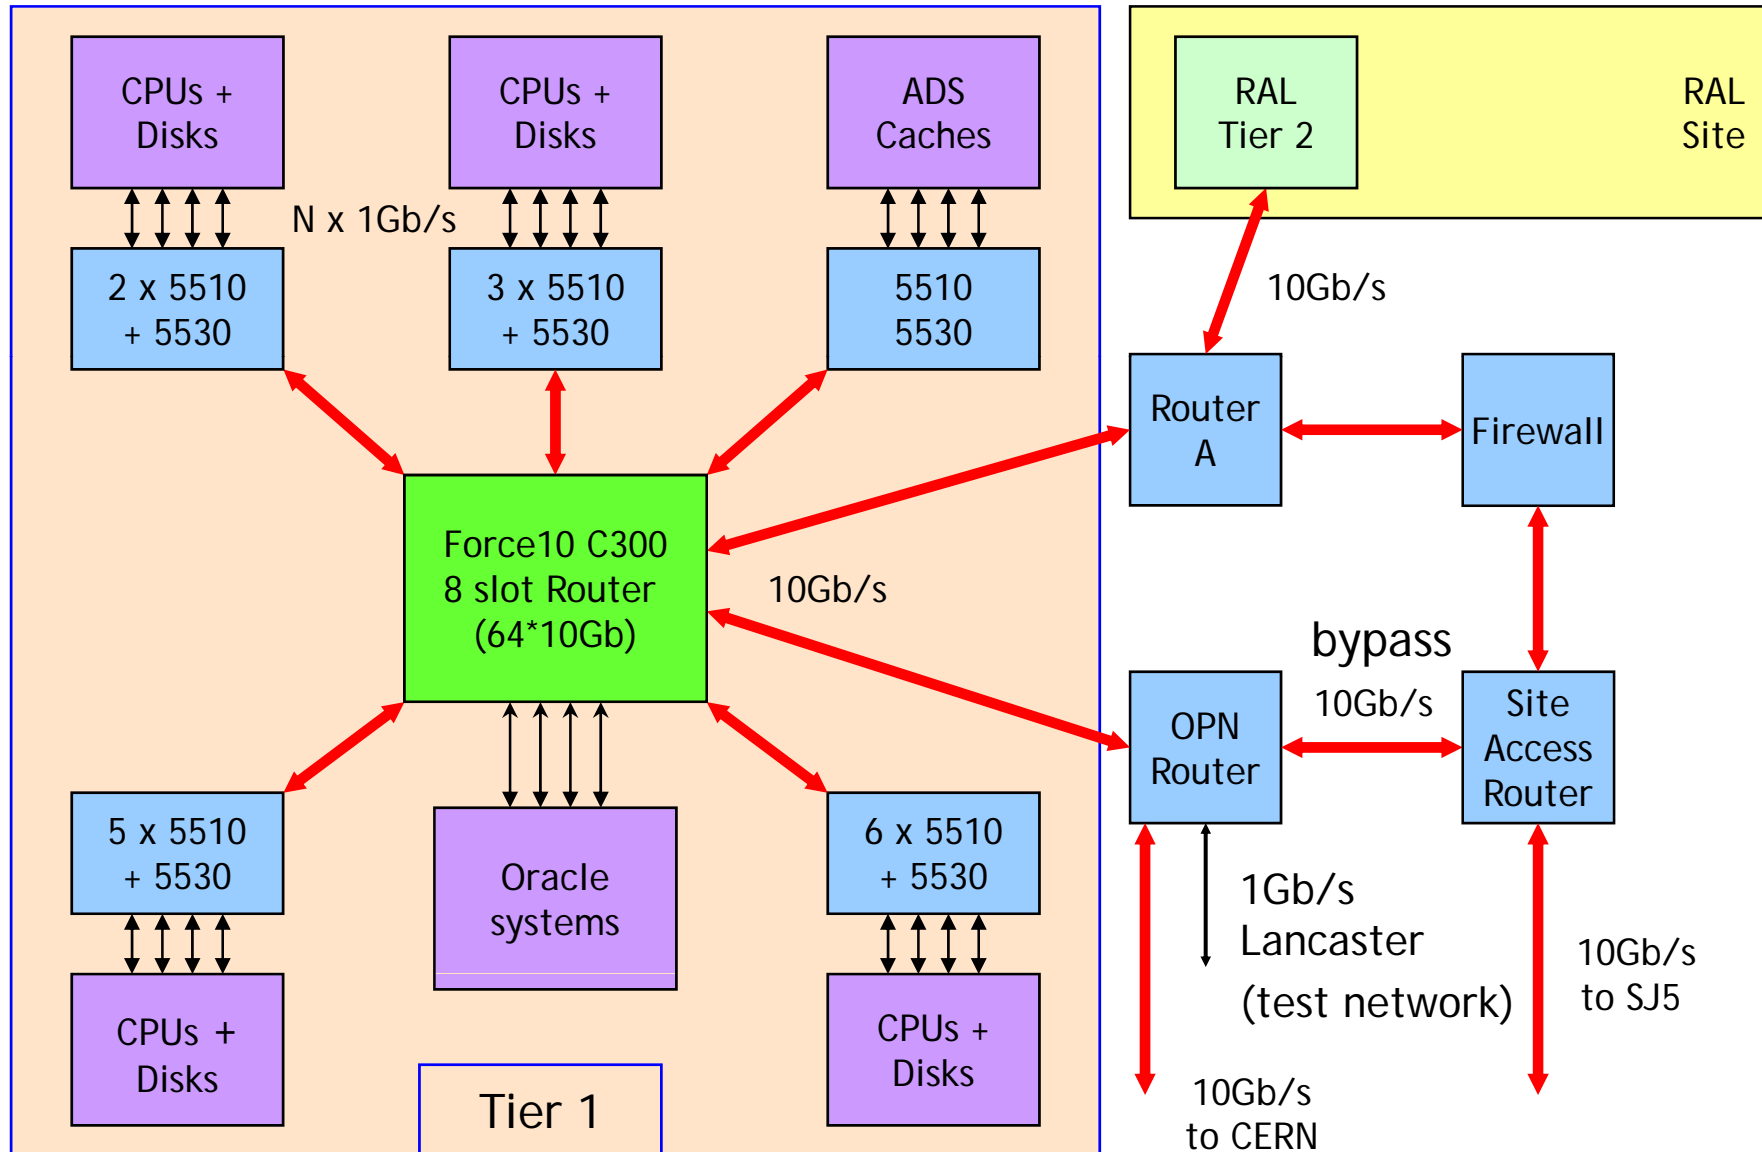
Hardware: CPU

- Production about 1500KSI2K on 600 systems
 - Recently upgraded about 50% of capacity to 2GB/core
- FY07/08 procurement (~3000KSI2K - but YMMV)
 - Streamline
 - 57 x 1U servers (114 systems, 3 racks), each system:
 - dual Intel E5410 (2.33GHz) quad-core CPUs
 - 2GB/core, 1 x 500GB HDD
 - Clustervision
 - 56 x 1U servers (112 systems, 4 racks), each system:
 - dual Intel E5440 (2.83GHz) quad-core CPUs
 - 2GB/core, 1 x 500GB HDD
 - Configuration based on 15kW per rack maximum, from supplied 'full-load' power consumption data

Hardware: non-Capacity

- Servers for Grid services (CEs, WMSs, FTS etc)
 - 11 'twin' systems, same as batch workers but two disks
- Low capacity storage
 - 6 x 2U servers, 8GB RAM, dual chip dual-core AMD CPU, 2 x 250GB HDD (RAID1 system), 4 x 750GB HDD (RAID 5 data), 3Ware controller
 - For AFS and Home filesystem, installation repositories...
- Xen
 - 4 'monster' systems for virtualisation
 - 2 x dual core AMD 2222 CPUs, 32GB RAM, 4 x 740 GB HDDs on HW RAID controller
 - For PPS service
- Oracle Databases
 - 5 x servers (redundant PSUs, HW RAID disks) and 7TB data array (HW RAID)
 - To provide additional RAC nodes for 3D services

Hardware: Network





Monitoring / On-Call

- Cacti - network traffic and power
- Ganglia - performance
- Nagios - alerts
- 24x7 callout now operational
 - Using Nagios to signal existing pager system to initiate callouts
 - Working well
 - Still learning
- Blogging
 - UK T2s and the Tier1 have blogs:
 - <http://planet.gridpp.ac.uk/>



Power Failure: Thursday 7th Feb ~12:15

- Work on building power supplies since December
 - Down to 1 transformer (of 2) for extended periods (weeks). Increased risk of disaster
 - Single transformer running at (close to) maximum operating load
 - No problems until work finished and casing being closed
 - control line crushed and power supply tripped!
- First power interruption for over 3 years
- Restart (Effort > 200 FTE hours)
 - Most Global/National/Tier-1 core systems up by Thursday evening
 - Most of CASTOR/dCache/NFS data services and part of batch up by Friday
 - Remaining batch on Saturday/Sunday
 - Still problems to iron out in CASTOR on Monday/Tuesday
- Lessons
 - Communication was prompt and sufficient but ad-hoc
 - Broadcast unavailable as RAL run the GOCDB (now fixed by caching)
 - Careful restart of disk servers slow and labour intensive (but worked) will not scale

<http://www.gridpp.rl.ac.uk/blog/2008/02/18/review-of-the-recent-power-failure/>