

GridPP

UK Computing for Particle Physics

Tier1 Site Report

HEPSysMan

30 June, 1 July 2011

Martin Bly, STFC-RAL



Science & Technology Facilities Council

e-Science

- RAL Stuff
- Building stuff
- Tier1 Stuff



- Email Addressing:

- Removal of old-style f.blogs@rl.ac.uk email addresses in favour of the cross-site standard fred.bloggs@stfc.ac.uk
 - (Significant resistance to this)
- No change in aim to remove old-style addresses but...
- ... mostly via natural wastage as staff leave or retire
- Staff can ask to have their old-style address terminated

- Exchange:

- Migration from Exchange 2003 to 2010 went successfully
 - Much more robust with automatic failover in several places
 - Mac users happy as Exchange 2010 works directly with Mac Mail so no need for Outlook clones
- Issue for exchange servers with MNLB and switch infrastructure
 - Providing load-balancing
 - Needed very precise instructions for set up to avoid significant network problems

- UPS problems
 - Leading power factor due to switch-mode PSUs in hardware
 - Causes 3KHz 'ringing' on current, all phases (61st harmonic)
 - Load is small (80kW) compared to capacity of UPS (480kVA)
 - Most kit stable but EMC AX4-5 FC arrays unpredictably detect supply failure and shut down arrays
 - Previous possible solutions abandoned in favour of:
 - **Local isolation transformers in feed from room distribution to in-rack distribution: Works! 😊**

- New structure within e-Science:
 - Castor Team moved into Data Services group under Dave Corney
 - Other Tier1 teams (Fabric, Services, Production) under Andrew Sansum
- Some staff changes:
 - James Thorne, Matt Hodges, Richard Hellier left
 - Jonathan Wheeler passed away
 - Derek Ross moved to SCT on secondment
- Recruiting replacements 😊


- Site
 - Sporadic packet loss in site core networking (few %)
 - Began in December, got steadily worse
 - Impact on connections to FTS control channels, LFC, other services
 - Data via LHCOPN not affected other than by control failures
 - Traced to traffic shaping rules used to limit bandwidth in firewall for site commercial tenants. These were being inherited by other network segments (unintentionally!)
 - Fixed by removing shaping rules and using a hardware bandwidth limiter
 - Currently a hardware issue in link between SAR and firewall causing packet loss
 - Hardware intervention Tuesday next week to fix
- LAN
 - Issue with a stack causing some ports to block access to some IP addresses: one of the stacking ports on the base switch faulty
 - Several failed 10GbE XFP transceivers

- Looking at structure of Tier1 network
 - Core with big chassis switches, or
 - Mesh with many top-of-rack switches
 - Want make use of 40GbE capability in new 10GbE switches
 - Move to have disk servers, virtualisation servers etc @ 10GbE as standard
- Site core network upgrades approved
 - New core structure with 100GbE backbones and 10/40GbE connectivity available
 - Planned for next few years

- Tier1 is a subnet of the RAL /16 network
- Two overlaid subnets: 130.246.176.0/21 and 130.246.216/21
- Third overlaid /22 subnet for Facilities Data Service
 - To be physically split later as traffic increases
- Monitoring: Cacti with weathermaps
- Site SJ5 link:
 - 20Gb/s + 20Gb/s failover
 - direct to SJ5 core
 - two routes (Reading, London)
- T1 <-> OPN link: 10Gb/s + 10Gb/s failover, two routes
- T1 <-> Core 10GbE
- T1 <-> SJ5 bypass: 10Gb/s
- T1 <-> PPD-T2: 10GbE
- Limited by line speeds and who else needs the bandwidth

RAL PSCS/Tier 1 Network Load

Updated: Jun 30 2011 11:10:33



- Summary of previous report(s):
 - 36 SuperMicro 4U 24-bay chassis with 2TB SATA HDD (10GbE)
 - 13 x SuperMicro Twin²: 2 x X5650, 4GB/core, 2 x 1T HDD
 - 13 x Dell C6100: 2 x X5650, 4GB/core, 2 x 1T HDD
 - Castor (Oracle) databases server refresh: 13 x Dell R610
 - Castor head nodes: 16 x Dell R410
 - Virtualisation: 6 x Dell R510, 12 x 300GB SAS, 24GB RAM, 2 x E5640
- New since November
 - 13 x Dell R610 tape servers (10GbE) for T10KC drives
 - 14 x T10KC tape drives
 - Arista 7124S 24-port 10GbE switch + twinax copper interconnects
 - 5 x Avaya 5650 switches + various 10/100/1000 switches

Tier1 Hardware Summary

- **Batch:**
 - ~65,000 HS06 from ~6300 cores
 - 750 systems (2,4,6 per chip, 2 chips/system)
 - 2, 3 or 4 GB RAM per core
 - Typically at least 50GB disk per core, some with two disks
 - 1GbE per system
- **Storage:**
 - ~8000TB in ~500+ servers
 - 6,9,18,38 or 40TB(ish) per server
 - 1GbE, 2010 generation on 10GbE
 - 10000-slot tape library, 500GB, 1TB or 5TB per cart
- **Network:**
 - Force10 C300 switch(es) in core
 - Stacks of Avaya (Nortel) 55xx and 56xx
 - Arista and Fujitsu 10GbE switches
- **Services:**
 - Mix of mostly old IBM and Transtec, mid-age SM twins and Transtec, and newer Dell systems

- Stats as of June 2011:
 - 16 million files: Used/Total capacities: 3.2PB/7PB on tape and 3.7PB/7.5PB on disk
- Recent news:
 - Major upgrade (2.1.9) during late 2010, which brought us:
 - Checksums for all files, xrootd support, proper integrated disk server draining
 - Minor upgrade (2.1.10-0) during February 2011 with bugfixes
 - Minor upgrade (2.1.10-1) next week, which brings us T10KC support
 - New (non-Tier1) production instance for Diamond synchrotron. Part of a new complete **Facilities Data Service** which provides data transparent aggregation (StorageD) metadata service (ICAT) and web frontend to access data (TopCAT)
- Coming up:
 - Move to new database hardware and better resilient architecture (using DataGuard) later this year for Tier-1 databases

- One of two batches of the FY09/10 capacity storage failed acceptance testing: 60/98 servers (~2.2PB) ☹
 - Cards swapped (LSI -> Adaptec)
 - Released for production use
- SL08 batch failing in production over extended period
 - Single-drive throws cause array lock up and crash (array + data loss)
 - Whole batch (50/110) rotated out of production (data migrated)
 - Updated Areca firmware, recreate arrays from scratch, new file systems etc
 - Subjected to aggressive acceptance tests
 - Passed, but...
 - ... issue with controller 'crashing ports for failed drives ~50% of cases
 - Seeking fix from Areca
 - Available for production in D0T1 service classes and accept some drive throws will need a reboot to see new disks

- Evaluated MS Hyper-V (inspired by CERN's successes) for services virtualization platform
 - Offers sophisticated management/failover etc without punitive cost of VMWare
 - However as Linux admins, sometimes hard to know if problems are due to ignorance of the MS world 😊
- Struggled for a long time with Infortrend iSCSI storage arrays (and poor support)
 - abandoned them recently and problems seem resolved
 - Evaluating Dell EqualLogic units on loan
- Have learnt a lot about administering Windows servers....
- Ready to implement production platform for Local storage HVs
 - 6 x Dell R510, 12 x 300GB SAS, 24GB RAM, 2 x E5640 for local storage HV
 - (14 x Dell R410, 4 x 1TB SATA, 24GB RAM for shared storage HV)

- Quattor
 - Batch and Storage systems under Quattor management
 - ~6200 cores, 700+ systems (batch), 500+ system (storage)
 - Significant time saving
 - Significant rollout on Grid services node types
- CernVM-FS
 - Major deployment at RAL to cope with software distribution issues
 - Details in talk by Ian Collier (next!)
- Databases
 - Students working on enhancements to the hardware database infrastructure

Questions?



30/06/2011

Tier1 Site Report - HEPsMan Summer 2011



Science & Technology Facilities Council
e-Science