

GridPP

UK Computing for Particle Physics

RAL Site Report

HEPiX Spring 2013, Bologna

15th - 19th April

Martin Bly, STFC-RAL



Science & Technology Facilities Council

e-Science



- **Site Power**

- Replacement of switch gear in power feeds to RAL campus now complete and redundant feeds reinstated

- **In-row cooling**

- Fitting of Double Regulated Valves to chilled water lines to in-row cooling completed in machine room.
- This will allow for better tuning of water flows to get more cooling from the in-row units without compromising the standard cooling regime

- Summary of new hardware:
 - 46 x ~90TB (34x3TB) storage servers, 10GbE, in two batches
 - One batch was upgraded to SAS drives
 - 14 x quad-system chassis in two batches
 - Each system has either E5-2660 or E5-2670 CPUs, 4GB RAM per thread, HT on, 2 x 1GbE teamed NIC, 2 x 2TB HDD (RAID0)
 - Noted that supplier benchmarking was not wholly accurate
 - Various Dell R3xx, R4xx, R5xx, R6xx servers
 - Various non-capacity needs including tape servers
 - Some with 10GbE NICs
 - More Force10 S4810P and S60 switches
 - Extreme x670V routers (2 off)
 - 40GbE Infrastructure cabling and patch leads

- Previously no Hyperthreading on WNs
 - Recent WNs have 3 or 4 GB RAM per core (no HT)
- Enabled progressively on batch service at end of 2012
 - Following period of extensive testing
 - Looking for load issues, power issues, ‘sweet spots’ for each generation
- Overcommit depends on node configuration
 - E.g. 2010 / 2011 generations of WNs:
 - 12 cores (24 threads)
 - 48GBytes RAM
 - Enable 20 jobs
- 2012/13 generation has 4GB/HT to enable 1 job per HT core

- Asymmetric Data Transfer rates in/out of Tier1
 - Solved - we hope!
 - Core C300 replaced by a S4810P
 - More ports including some 40GbE
 - More recent firmware
 - 10Gb/s choke on Tier1 outbound link to WAN
 - When the network is not overloaded, the rates and latencies are much better
- LAN
 - Planning well advanced for Tier1 migration to Mesh network
 - Final push to ensure the WAN transfers rates are the best possible will be done after the changeover

Service	Version	VM?
APEL	UMD-2 / SL5	Yes
ARGUS	EMI-2 / SL6	Yes
BDII (site)	EMI-2 / SL5	Yes
BDII (top)	EMI-2 / SL5	No
CREAM CE	EMI-2 / SL6	Yes
FTS	EMI-1 / SL5 (& FTS3 Test)	Yes
MyProxy	UMD-2 / SL5	Yes
L&B	EMI-2 / SL5	Yes
WN	EMI-2 / SL5 (Some SL6 tests)	No
LFC	EMI-2 / SL5	Yes
Torque server	UMD-1 / SL5	No
WMS	EMI-2 / SL5	Yes
UI	EMI-2 / SL5 & UMD-1 SL5	No

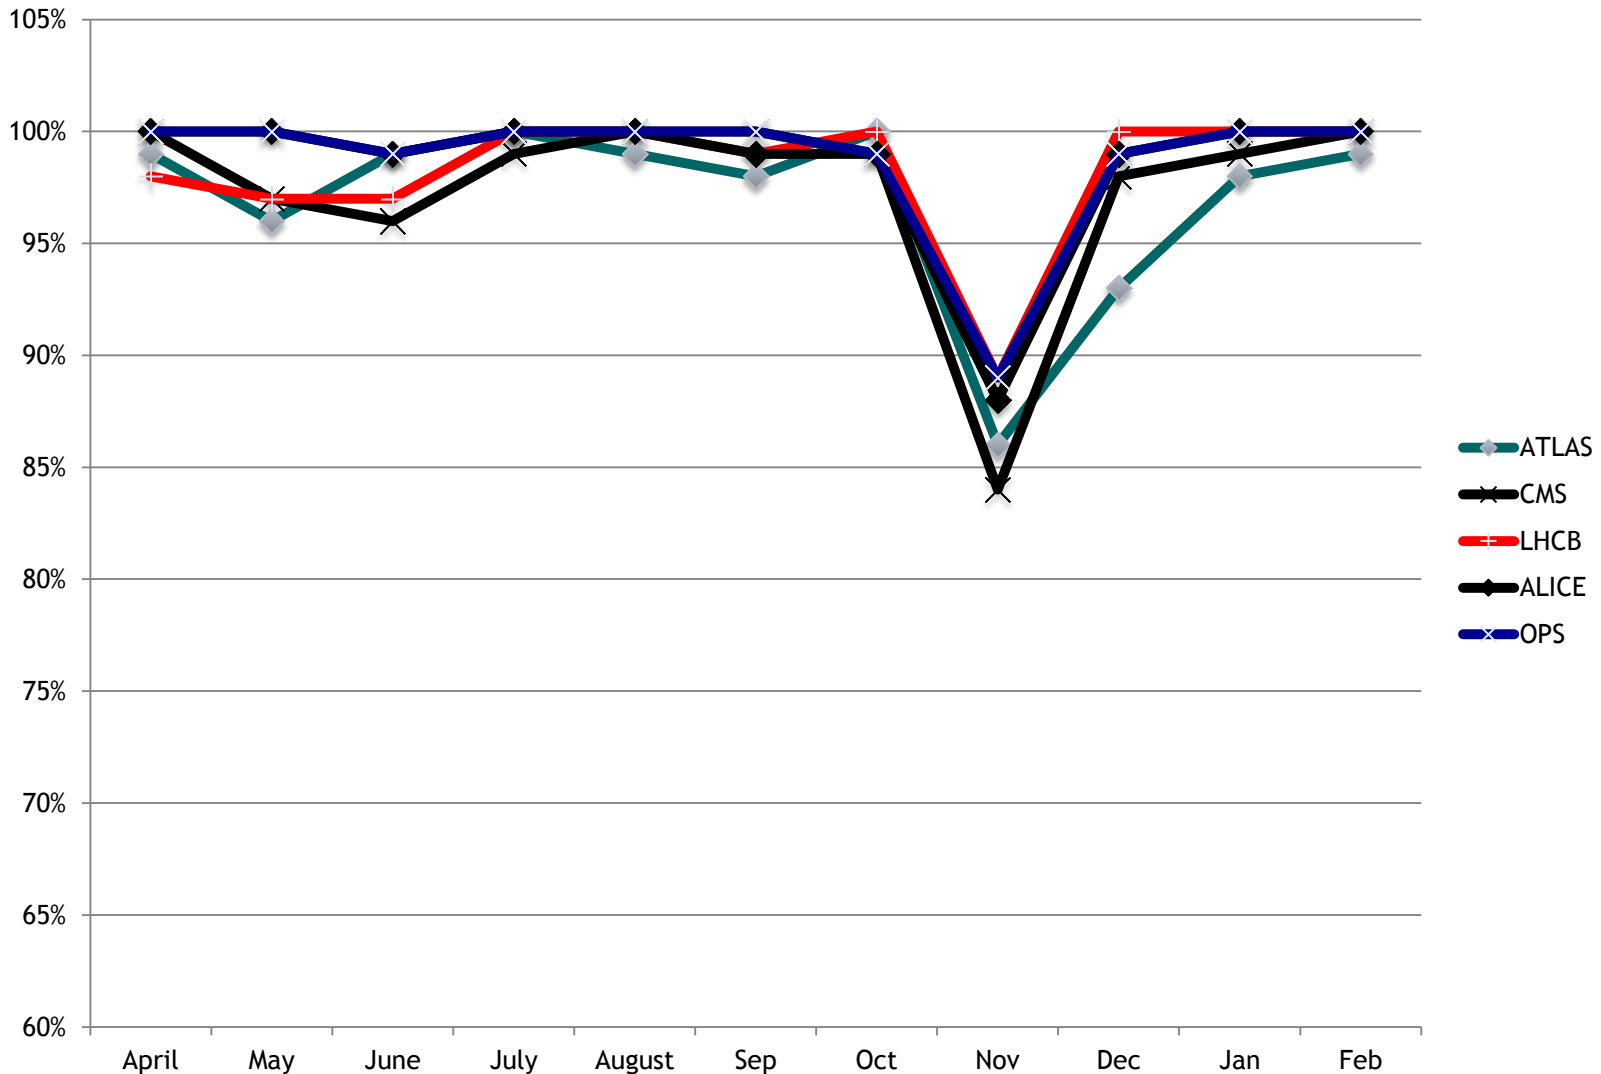
- **CASTOR now running at version 2.1.12**
 - Database upgraded to Oracle 11.2
 - Replaced LSF scheduler with CERN transfermanager
- **Stability significantly improved**
 - A few persistent timeout failures remain, mainly effecting CMS
- **But disk server draining remains slow and error-prone**
 - Needs DB intervention
- **Working on next upgrade (2.1.13)**
 - Improved logging and better disk server draining tools

- RAL site network stability
 - Load-related issues on main internal router and firewall
 - Multiple causes - no simple/single answer
 - One incident caused by xrootd fallback from Tier-2 to Pisa.
 - Previous incident was a host initiating an outbound scan
 - Solution - upgrades and simplification/restructuring
- Batch server start rate
 - Torque/Maui server cannot start short jobs fast enough
 - Looking at alternatives but problem manageable
- Long setup times causing ATLAS/LHCb job timeouts
 - CVMFS?

- And then of course, we had two power supply problems...
- ...within two weeks:
- 7th November: Site power supply failure
 - ‘Caused’ by routine switching operation at a leaf switch panel generating a fault which tripped the main transformer...
 - ... while we were only operating on one...
 - ... and the UPS generator failed to take the load
- 20th November: Power surge
 - Due to error in bypass feed work for UPS feed
 - Extensive damage to equipment
- See my later talk @ this meeting...



Tier-1 Availability Apr12 - Feb13



- Next generation (disk) storage project continues. Challenge: Try to find something simple and easy to run in conjunction with CASTOR for tape access.
 - Some interesting candidates (HDFS and Ceph) and an old stalwart (dCache) remain in the mix. Testing continues.
 - No clear obvious choice at the moment that can currently improve on just running CASTOR but interesting possibilities.
 - If we are to deploy something in time for run 2, must be soon. Consider options at Tier-1 review in May.
- Major outage (2 day) for electrical safety checks
- Essential power board upgrade
- FTS3, EMI-3, CASTOR 2.1.13, SL6
- Major Tier-1 network upgrade (May?)



Infrastructure as before:

- Nagios - 1642 hosts, ~41000 tests
- Ganglia - ~400,000 RRD files
- RT - for ticket tracking.
- Site Dashboard at: <http://www.gridpp.rl.ac.uk/status/>
- Twitter @raltier1

Four incidents investigated in 2012

One so far in 2013.

- [20120316 Network Packet Storm](#)
- [20120613 Oracle11 Update Failure](#)
- [20121107 Site Wide Power Failure](#)
- [20121120 UPS Over Voltage](#)
- [20130219 Disk Server Failure File Loss](#)