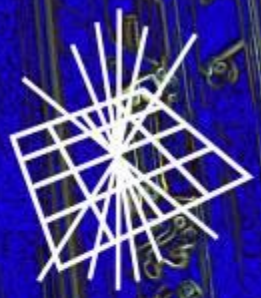

GridPP

UK Computing for Particle Physics

Oxford Site Report
HEPSYSMAN

Pete Gronbech
GridPP Project Manager

June 22nd 2016

Oxford Local Cluster

Current capacity
6000HS06 680TB

- Almost exclusively SL6 now.
- New procurement summer 2015.
- 4 Supermicro twin squared CPU boxes provide 256CPU physical cores
- Chose Intel E5-2630v3's should provide approx. ~4400 HS06 upgrade.

- Storage from Lenovo.
  - 1 U server with two disk shelves, containing 12 * 4TB SAS disks.
  - Provided an increased capacity of ~350TB
  - ~88TB for NFS and the rest for Lustre

GridPP
UK Computing for Particle Physics

Oxford's Grid Cluster

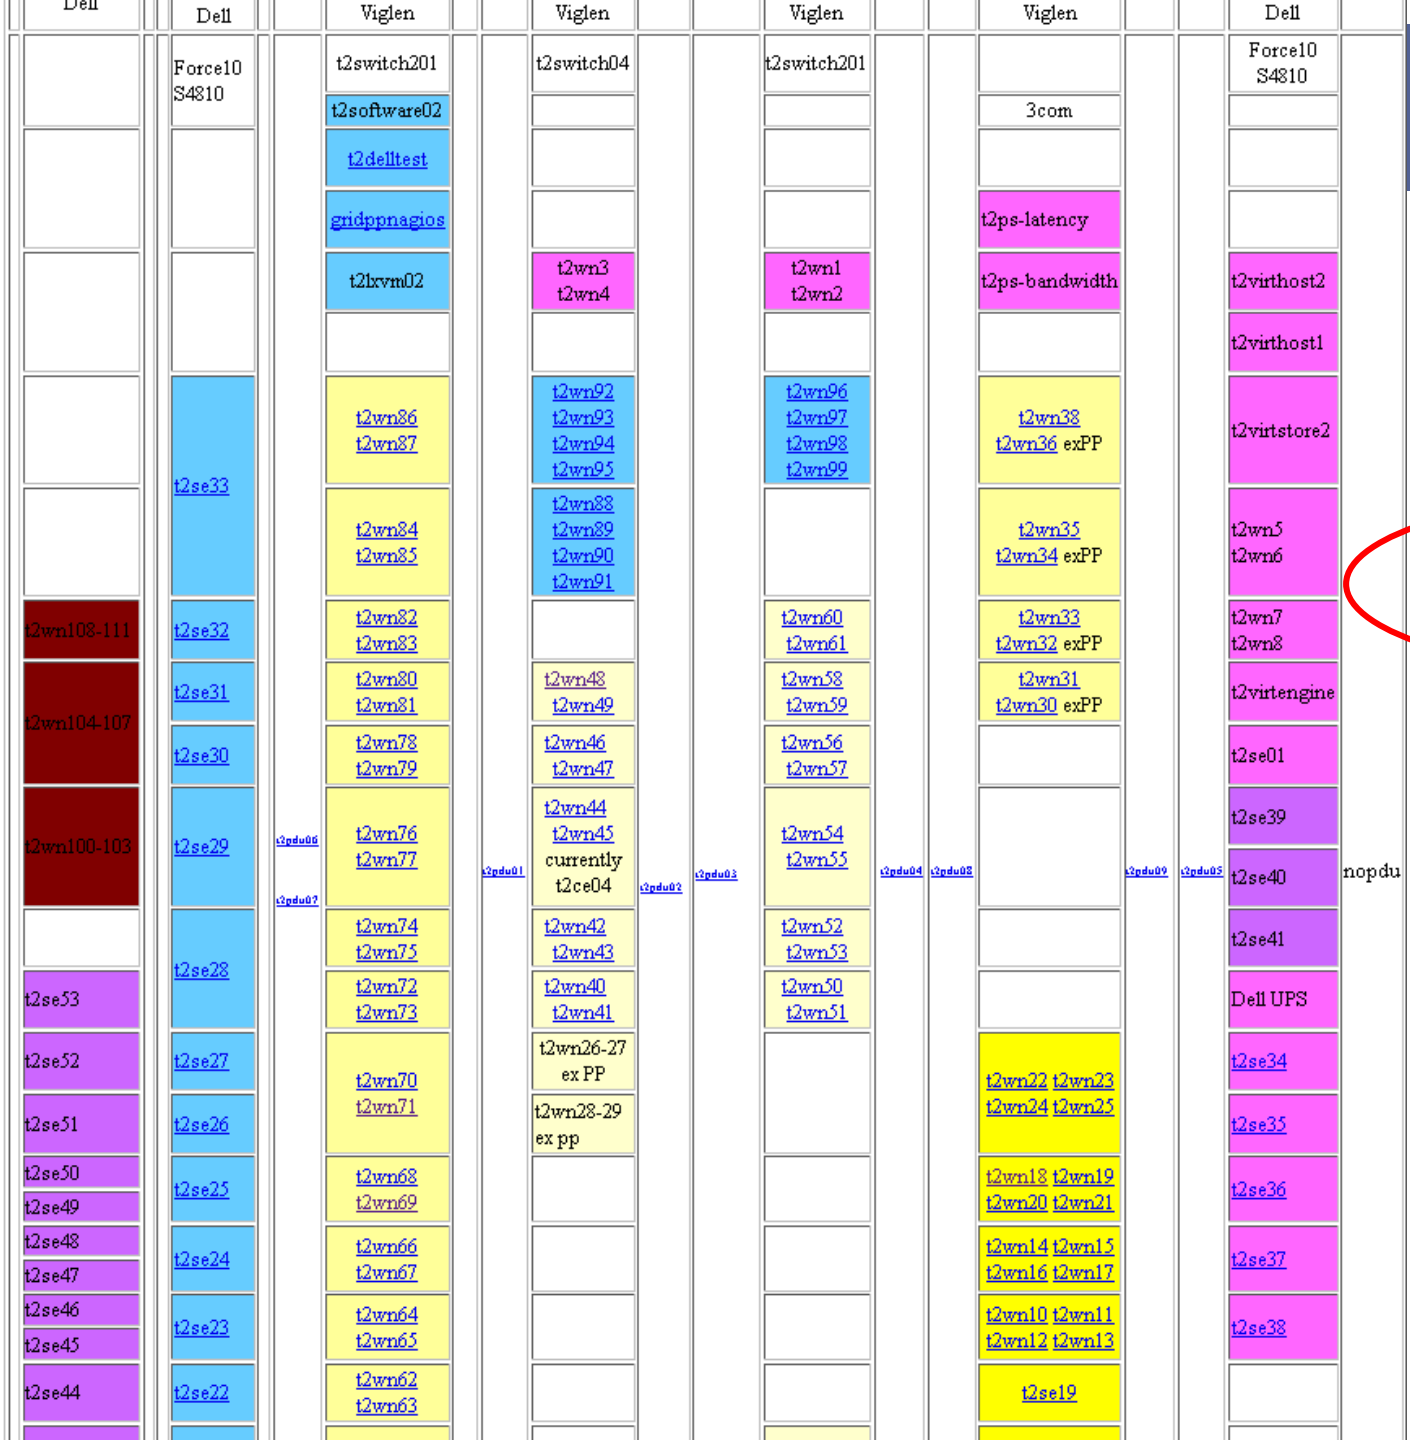

Supplier	Date	Warranty expires	WN CPU
Dell	2004	2007	
Viglen	Aug2007	2010	Intel 5345
Viglen	Sep 2008	2011	Intel 5420
Viglen	Nov 2010	2014	AMD 6128
Dell	May 2011	2016	Intel 5650
Dell	Q1 2012	2017	AMD 6276
Dell	Jan 2014	2019	
Viglen	Mar 2014	2019	Intel 2650v2

Oxford Grid Cluster

GridPP4 status
Autumn 2015

Current capacity
16,768HS06 980TB

- DPM Storage older Supermicro 26, 24 and 36 bay servers decommissioned, so capacity is provided by Dell 510 and 710s. 12 bay with 2 or 4 TB disks.
- Majority of CPU nodes are 'twin-squared' Viglen Supermicro worker nodes have been installed. Intel E5- 8 core (16 Hyper-threaded cores each) provides 1300 job slots with 2GB RAM.
- The majority of the Grid Cluster runs HT Condor behind an ARC CE.
- (Some viab nodes ~88 cores)

Decommissioning old storage

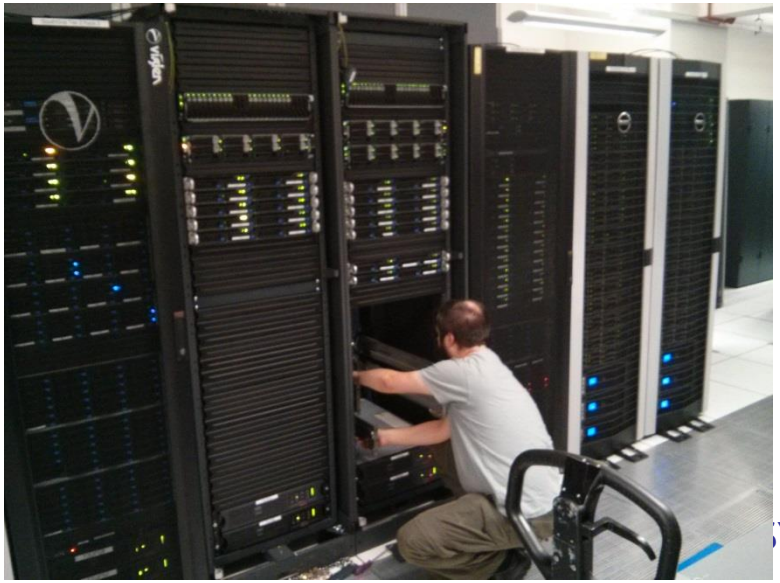

17 Supermicro servers removed from the DPM SE.

(Reduction of 320TB new total 980TB)

Storage March 2016

Old servers removed, switched off or repurposed

Software up to date

OS up to date

Simplified ✓

Plan for GridPP5

- CPU
  - GridPP4+ h/w money spent on CPU. This rack is identical to the kit used by Oxford Advanced Research Computing.
    - Initial plan is to plug into our Condor Batch as WNs
    - When staff levels and time allows, we will investigate integrating the rack into the ARC cluster.

CPU Upgrade March 2016

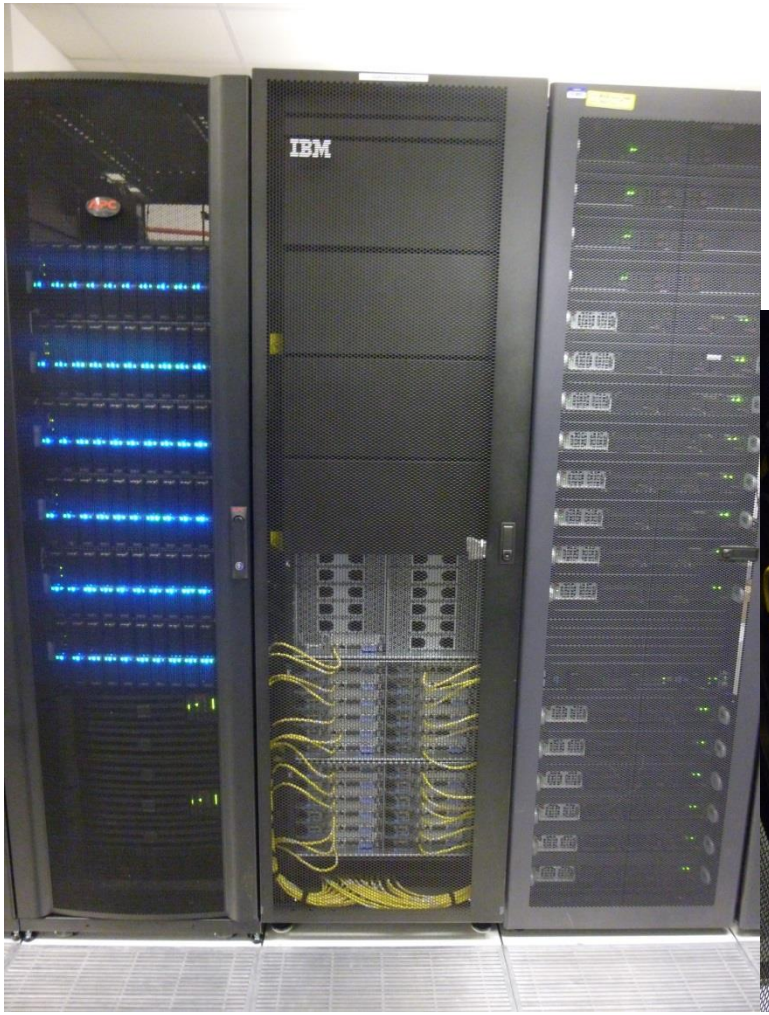

Lenovo NeXtScale

25 Nodes each with Dual E5-2640 v3 & 64GB RAM

800 new cores (new total ~2200)

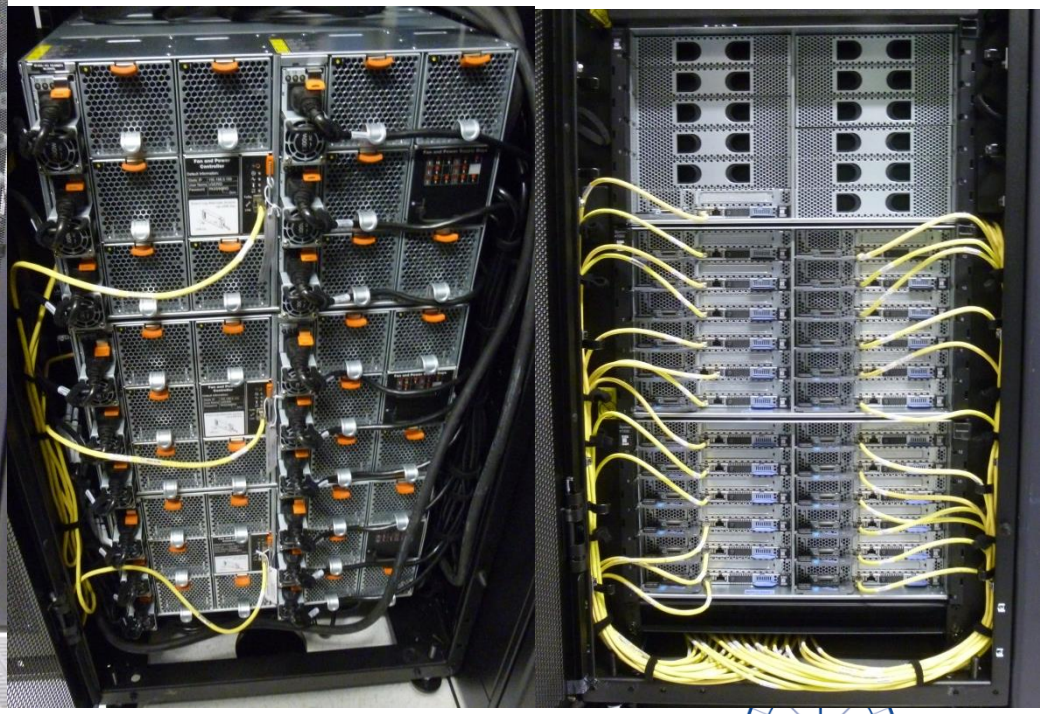

Begbroke Computer Room houses the Oxford Tier 2 Cluster

Spot the difference

Air Con Failure

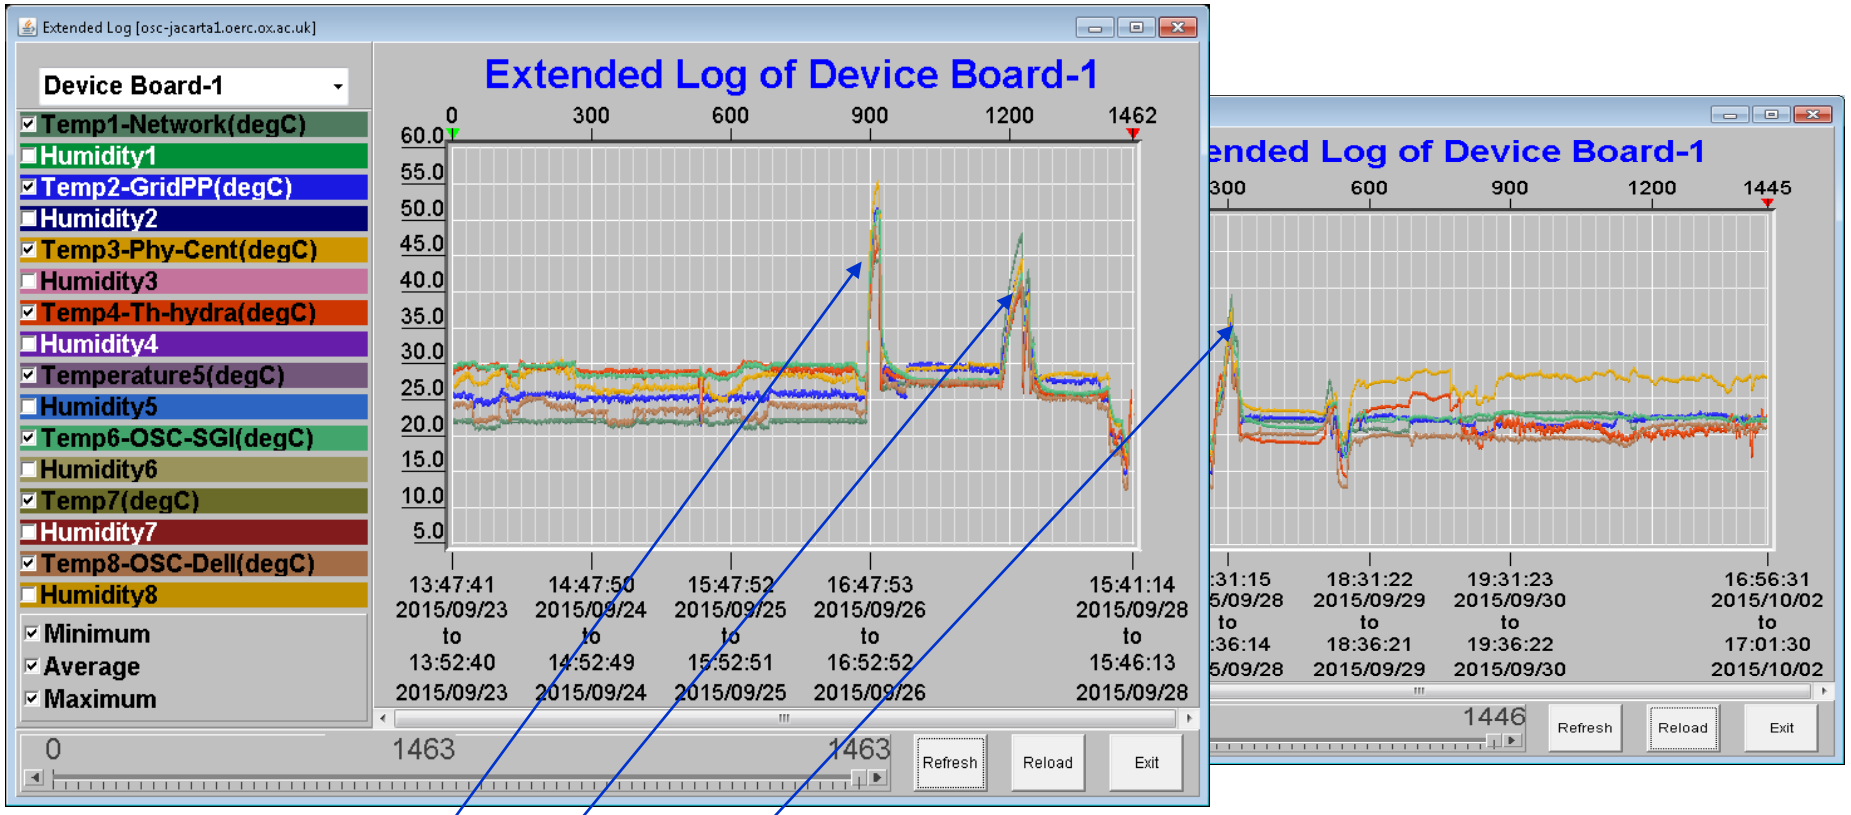

Saturday 16:19 Sept 26th 2015

Sunday 16:13

Monday 17:04

- People came into site or remotely switched off clusters very quickly.
- Building Services reset the A/C in 1-2 hours on both weekend days.

- The bulk of the load comes from University and Physics HPC clusters, but it turns out some critical University Financial services were also being run on this site.
- The incident was taken seriously and backup system ordered on Tuesday morning and installed from 10pm to 2am that night.
- Provides ~100KW back up in case of further trips. Normal load is ~230KW so main clusters were restricted.

Additional temporary A/C

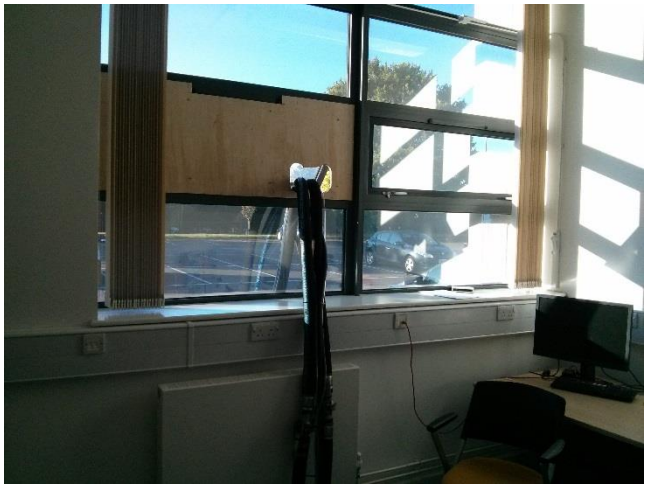

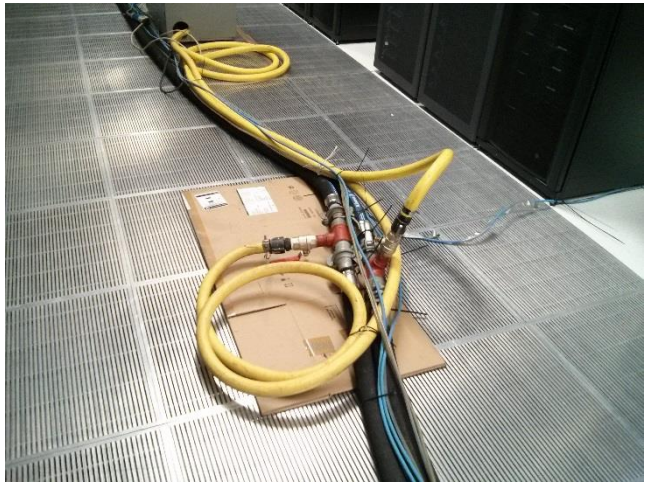

- Pressurization unit had two faults, both repaired. New improved unit to be installed on Monday.

- A 200KW computing load will heat up very very quickly when the A/C fails.
- It always seems to do this out of hours. You need to react to this very quickly.
- Really needs to be automated, even fast response from staff not fast enough.
- Even risk mitigation has it's own risks.

Water and Electrics!!

Here we go again

- Unfortunately a thunder storm on Sunday 12th June 2016 caused a large brown out at the Begbroke site.
- This caused some pumps to switch off.
- Pattern was very similar to last time, Physics and ARC staff alerted first by our monitoring.
- Too long before building services got to site.
- Still work to be done to improve this sort of call out.

Conclusions

- Local Cluster
  - 2015 Local Cluster upgrade went well now need to plan and purchase 2016 upgrade.
  - Recruitment on going.
- Grid Cluster
  - A time of streamlining and rationalisation.
  - Still a lot of work to do, to investigate integrating with university resources.
  - Will this be possible, will it save time, or allow bursting to greater resources?
  - Possibly benefits of cost savings on h/w maintenance and electricity costs.
- A/C problems
  - Need faster response from Building Services
  - Need auto shutdown of systems