

Edinburgh (ECDF) Site update

Wahid Bhimji
Andy Washbrook

And others including ECDF systems team

**Not a comprehensive update but what ever occurred to
me yesterday**

Edinburgh Setup Reminder

ECDF Grid Tier2

- ◉ Cluster shared with other uni users:
 - ~3000 cores in total
- ◉ Storage just for gridPP :
 - ~175 TB in DPM form on dense DELL boxes
- ◉ Systems team who do most hardware /os

Local computing

- ◉ Run physics-wide
 - Andy has to interact but I'm not going to talk much about hardware etc. much here
 - Will talk a bit about ATLAS user running on it

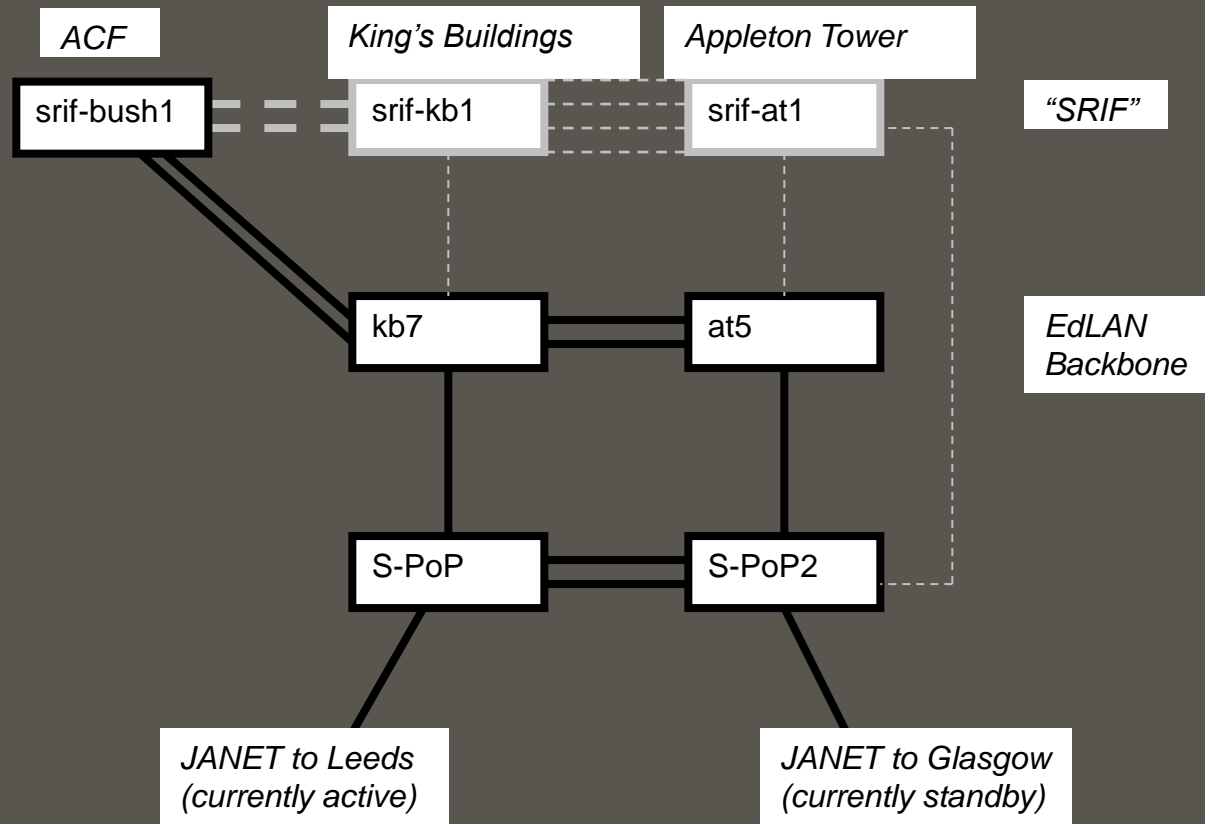
ECDF: New Kit

◉ GridPP hardware cash:

- most spent on storage
- (DELL R510 + 2 MD1200) *3 + 2 R510s
- 175 -> 350 (+) delivered TB (by june)
- Also middleware servers – soon all on newish hw

◉ DRI cash

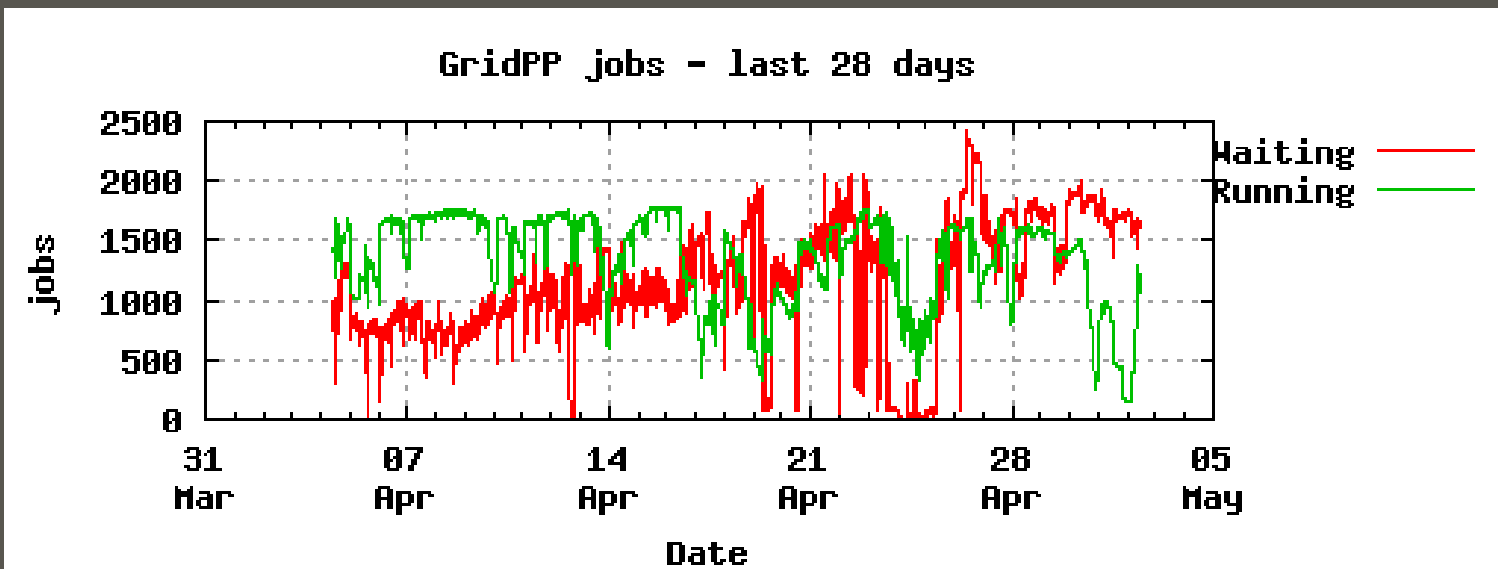
- mostly spent by networking guys improving outgoing links to 10 Gig resilient links
- (matching investment from uni)
- Also 10gig switch for pool servers; new racks etc.



New network layout – old links grayed out

Operations

- Generally running smoothly
- Bit of deployment: cvmfs; cream ces etc.
- Lots (for us) of running jobs (up to ~1700)
- Though should be able to run 3000

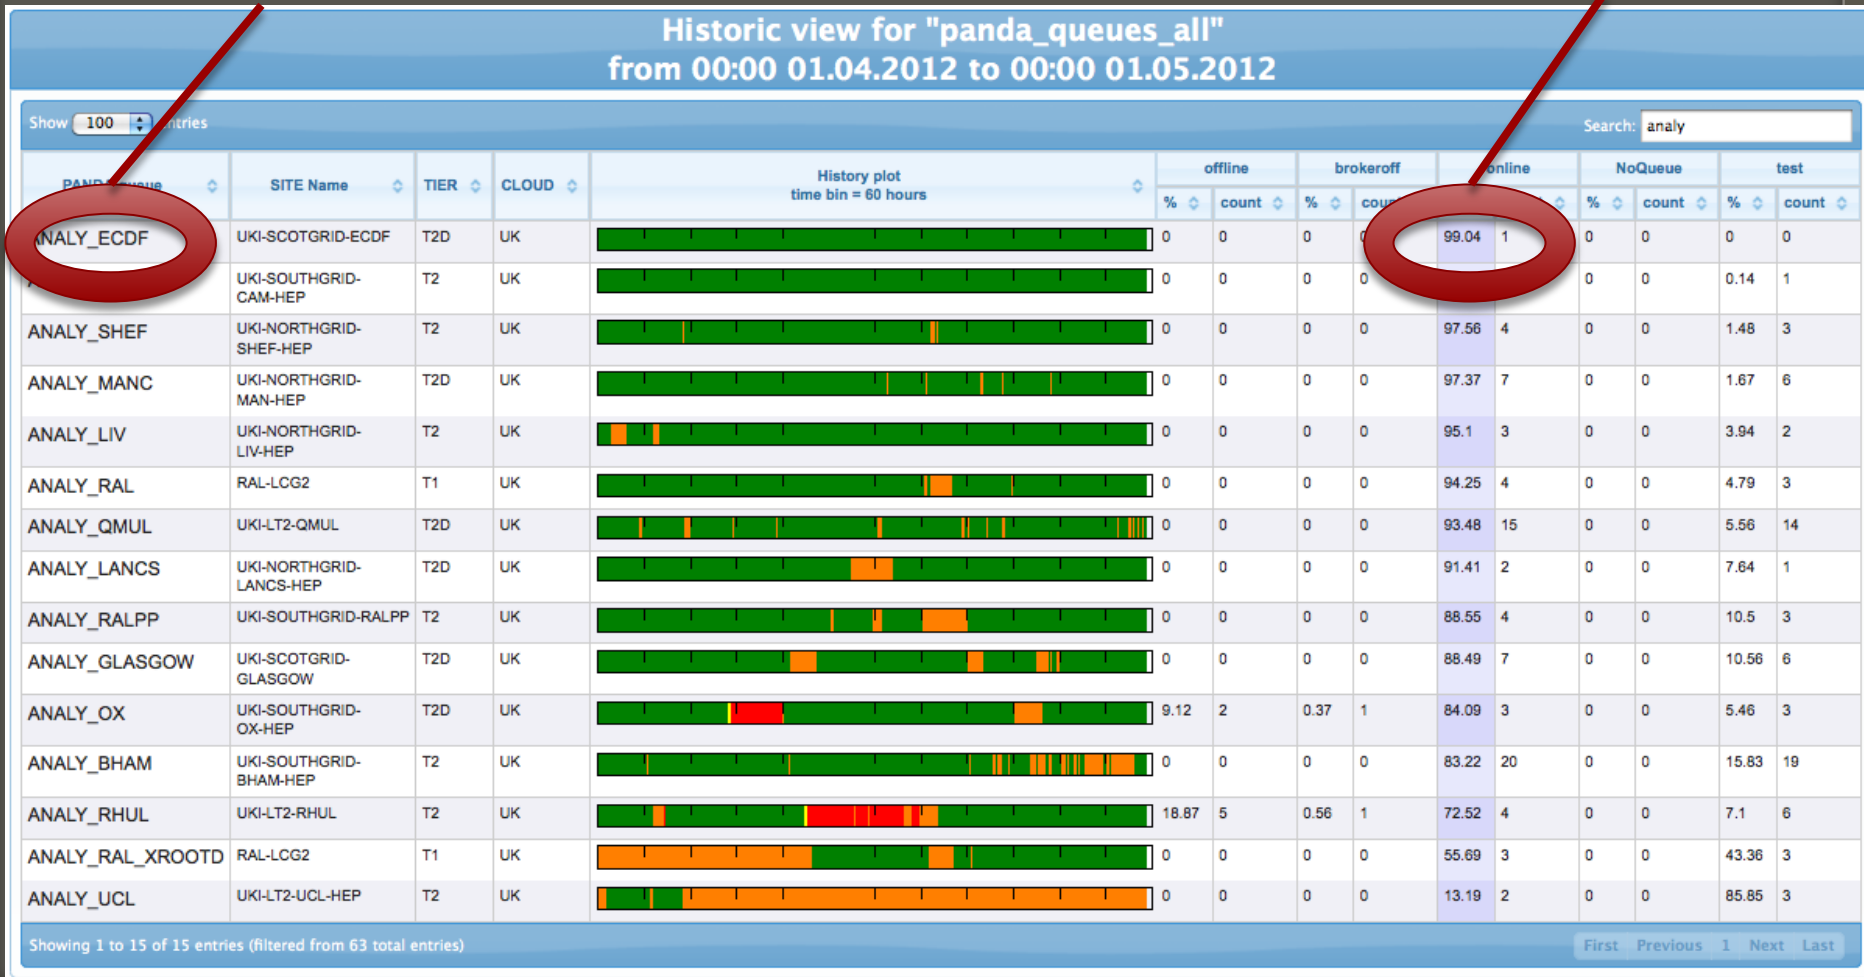


Availability: was great for ages...

ECDF

APRIL ATLAS Availability

99%



but all good things come to an end?

- **In last week qs have been up and down**
 - Cmtsite timeout
 - Other site see it – but seems to turn us off more?; and why now? Discuss?
- **Also datadisk blacklisting**
 - “alpha” site for ATLAS with decent no. of jobs means ATLAS expect to be able to place data
 - We don’t actually have bigboy level disk ...
 - but do have some being commissioned

Multicore

Involved in three areas for ATLAS multicore readiness:

- **AthenaMP performance measurement**

- CPU/ memory usage, serialisation timing, optimum job length and number of input files

- **ATLAS queue validation**

- Working with ATLAS distributed computing team to validate queue readiness for all sites wishing to run ATLAS multicore job
 - In UK: RAL, ECDF, Glasgow, Lancaster

- **Multicore scheduler simulation**

- Developed testbed to simulate scheduler response to jobs requesting different numbers of CPU cores
- Used to identify resource contention that leads to loss of job throughput

See Andy's talk at GridPP – and aCHEP 2012 talk:

"Multi-core job submission and grid resource scheduling for ATLAS AthenaMP"

Local computing: ATLAS Running on T2 files

- ATLAS users getting files locally with dq2
 - Get failures and running out of local disk space
 - So read interactively on desktop from T2 LOCALGROUPDISK instead
 - Transfer in managed with atlas central Datri tool
 - Files opened with rfio directly
 - They can run on ecdf batch (or condor) to scale up
 - Need few instructions for users but works fine
- https://www.gridpp.ac.uk/wiki/RFIO_Local_Access
- Users need to use ROOT TTreeCache for decent network read
 - Also Tree->GetEntriesFast to avoid slow jobs start if lots of small files
 - Obviously depends on network to T2.

Conclusions

ECDF running well (x-fingers)

- Has allowed time for some other interesting stuff
- We are getting quite a big site from ATLAS side:
 - T2D
 - Analysis jobs
 - GROUP space for SOFT-SIMUL
 - Multicore jobs
- That brings responsibility and expectations on hardware