Imperial College London

CMS Computing Resource Planning, going from 2012 to 2015

David Colling

Imperial College London

Apology and thanks

The Apology

This talk is rather dry, but it does have some important implications…

The Thanks

The figures I use are taken from a variety of people who do the work. Thanks to them all.

Imperial College London

Content

- What happened in the 2012 running
- What would 2015 look like if we didn't upgrade the computing model from the 2012 version
- How is the model evolving to help resources squeeze.
  - RECO becoming transient
  - Data Federations (AAA)
  - Blurring between different Tiers
  - CMSooooCloud (using the HLT)
  - Opportunistic Resources (speculative)

Will try to keep things conceptual for detailed numbers see the document that went to the C-RRB:

https://cms-docdb.cern.ch/cgi-bin/DocDB/RetrieveFile?docid=12055&version=5&filename=CMSResourceRequest_2013_FallV3.doc

and the future planning from the project office at:

https://cms-docdb.cern.ch/cgi-bin/DocDB/ShowDocument?docid=12170

2012 – a reminder

- Took data at ~1KHz (~400Hz prompt & ~600Hz parked)
- Took data for 25% longer than we expected.
- Software was released early, at the time of ICHEP, and stayed with it through the year. Prompt reco used for most analysis with only limited re-reco of specific samples. This was NOT the case in 2011.
- Pile up increased with instantaneous luminosity
- PDS overlap factor =1.25

2012 – a reminder

- Tier 0 (just about) kept up with the prompt reco.
  - Would fall behind when there were short gaps between fills and catch up in longer gaps between fills. Overall T0 “keep up factor”=0.5
  - A system for spilling T0 workflows into public queues was commissioned and heavily used.
  - The load at the T0 was higher than expected because of the need to repack the parked data for distribution to the Tier 1s
  - The CAF was under utilised, in fact, half the processing capacity was moved to the T0
  - Although when not busy did contribute re-reco

Imperial College London

2012 – a reminder

How does reco time vary with conditions?

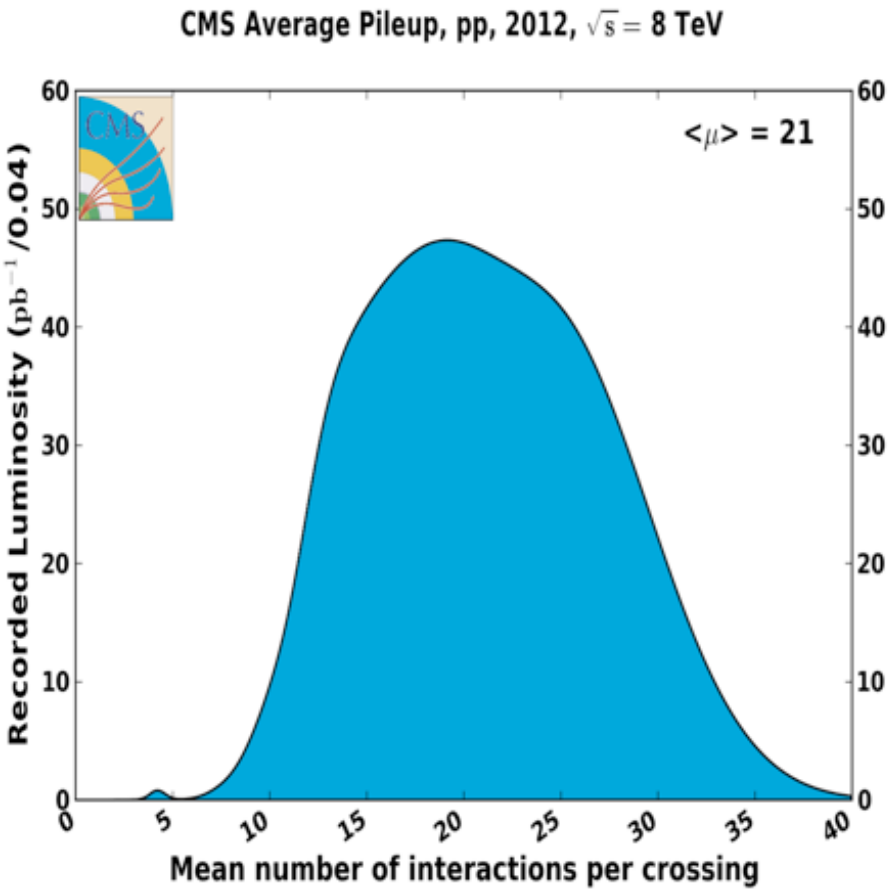

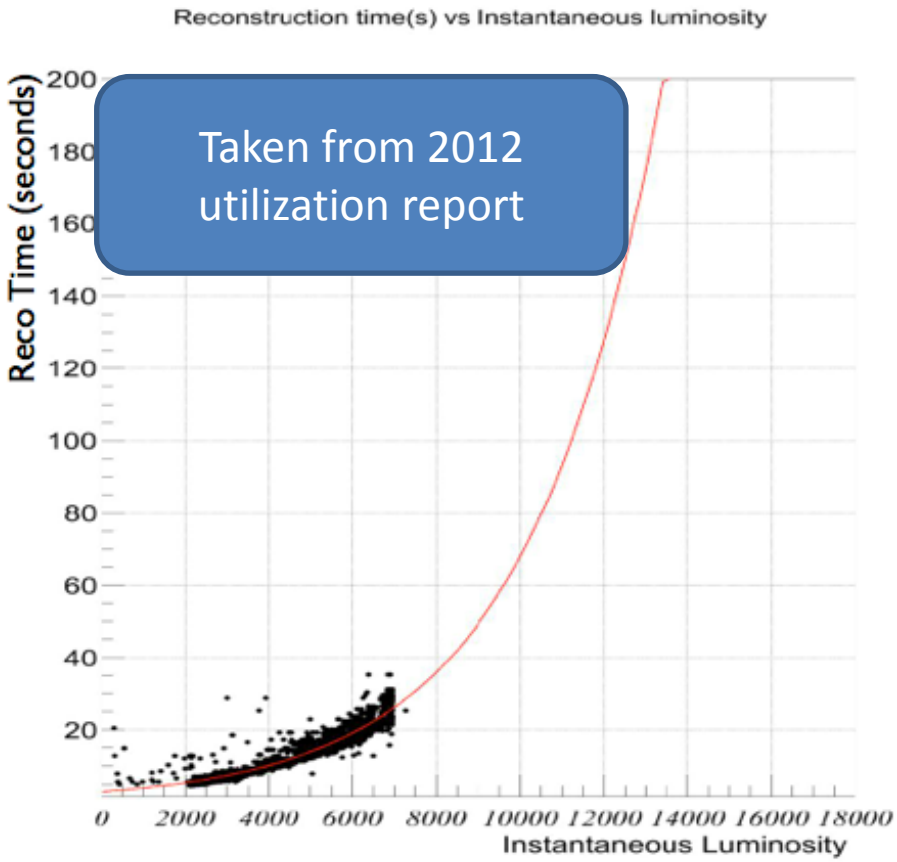

Taken from 2012 utilization report

Imperial College London

2012 – a reminder

- Tier 1 utilisation was much as expected (with some minor exceptions not worth mentioning)

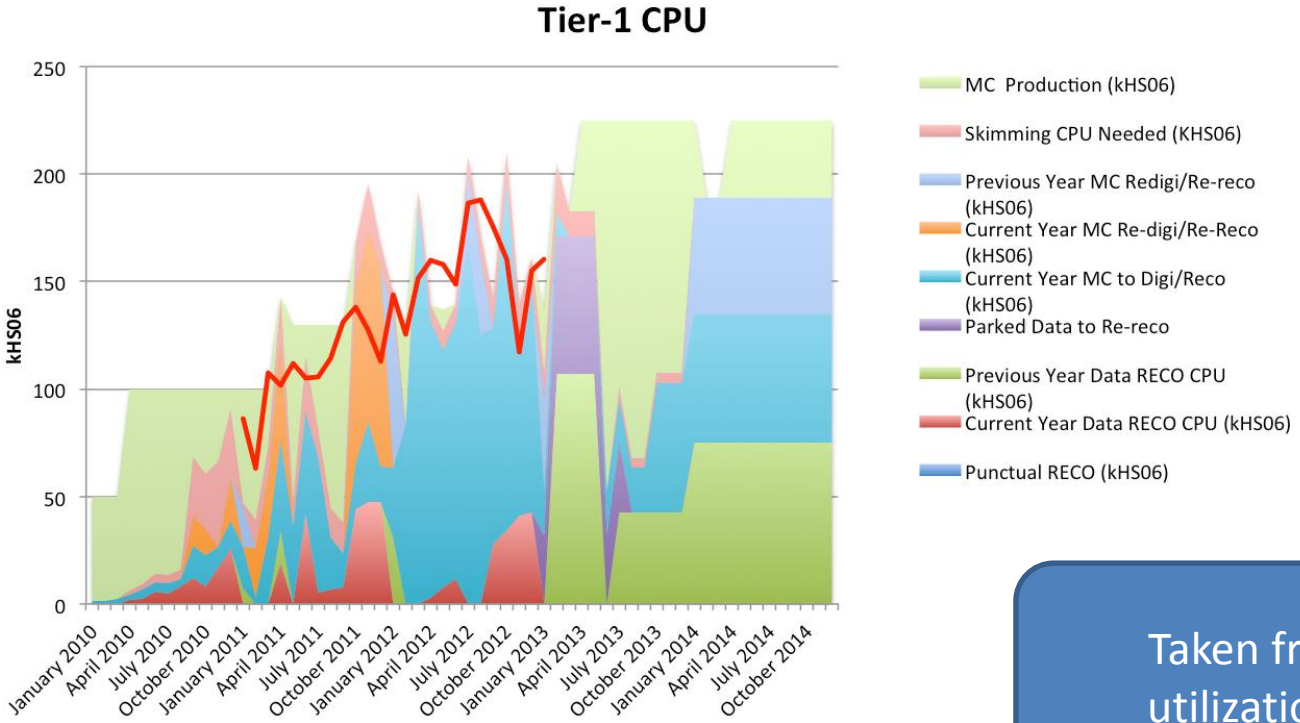

Taken from 2012 utilization report

Imperial College London

2012 – a reminder

- Tier 2s (where all the physics happens) only worked because countries delivered far more than was pledged.

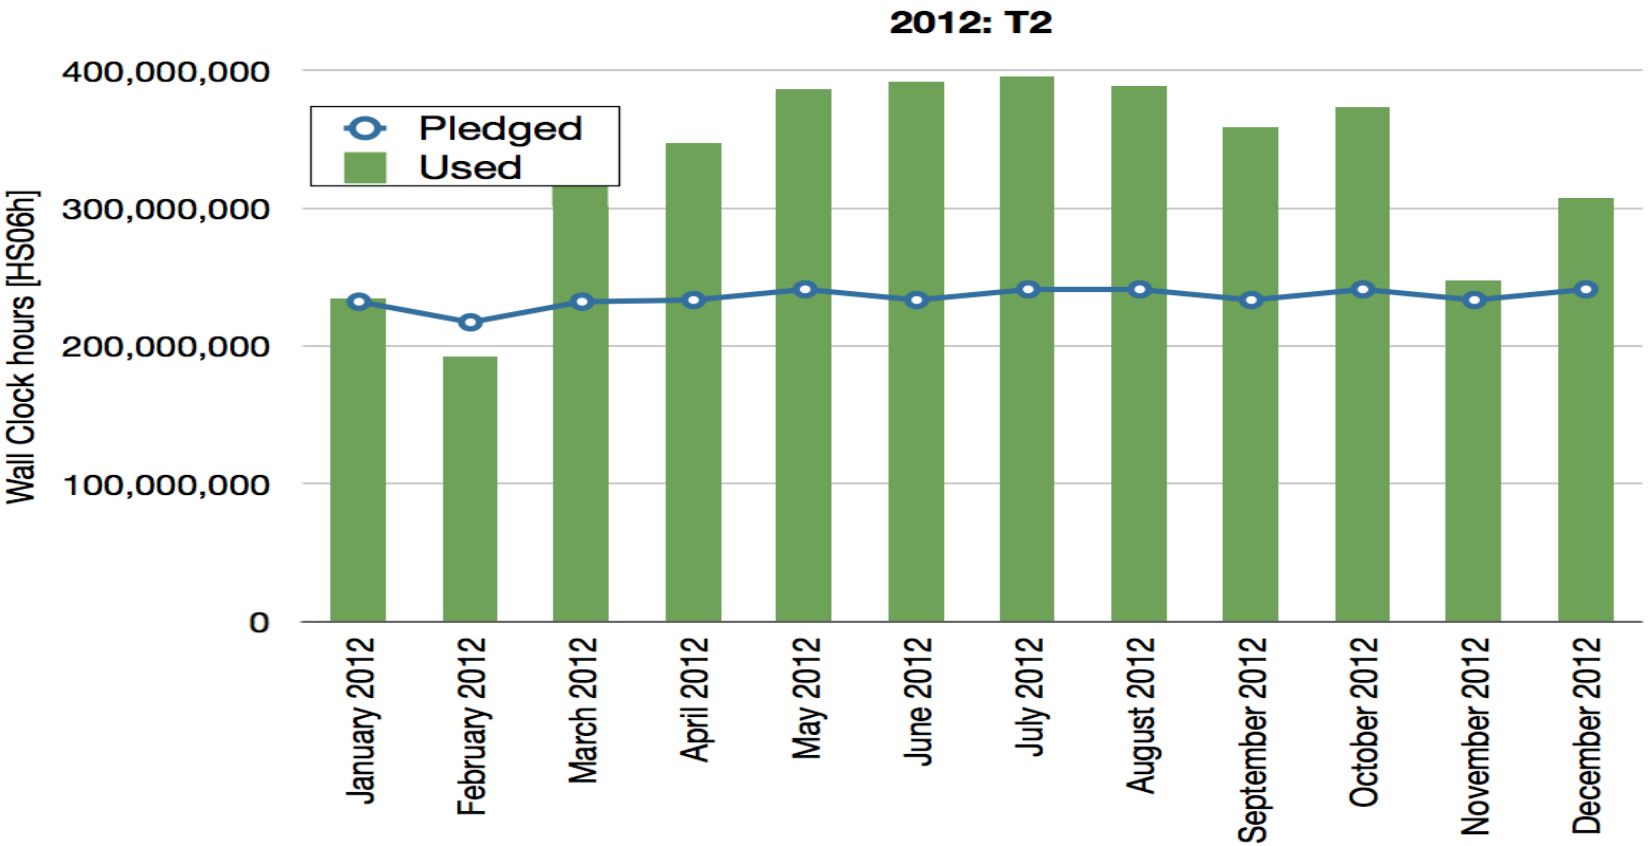

Going from 2012 to 2015

- No parked data in 2015 – no long shutdown in which to reprocess – but still desire to take data at (at least) 1KHz. Factor of 2.5 in timing
- Expected increased pile up scenario Factor 2.5 in timing
- Moving from 50ns to 25ns also initially gave a factor of 2, however offline have recovered this.
- Assume PDS overlap = 1.25
- Assume T0 keep up factor =0.4
- Assume 2012 (not 2011) level of re-reco organisation

Overall factor of 6 if we were to run as we did in 2012

This is all expressed in reco timing but equally visible in storage

So what are going to do about it?

- Make the RECO transitory
  - The original version of the CMS computing model assumes that one copy of the current version of the reconstructed data and the current year's raw data is kept on disk, with 10% of 2 copies of preceding versions of the RECO. In 2015 RECO will be transitioned to a transient format with only a sliding window of 3-6 months after processing retained on disk and not written to tape at all. This had allowed CMS to limit the new disk space needed at Tier-1s.

So what are going to do about it?

Traditional CMS Grid Model

- CMS generally moves jobs to data so that the job runs in the same room as the data are stored
  - Design decision from 10 years ago, based on assumption that data transfers are slow and unreliable
- Very successful, but experience has indicated problems
  - Longer queuing times than needed when free CPU not near data
  - Hard to incorporate resources not dedicated to experiment
  - Larger storage requirements than affordable in the future
  - Users prefer to run locally if possible, but might not have the data
- All of this can be solved if we could provide access to any data, anytime, anywhere!

Any Data, Anytime, Anywhere(AAA)

- Goal: make all data even more straightforwardly available to any CMS physicist, anywhere
  - Reliably: no access failures
  - Transparently: never notice where the data actually reside
  - Easily: no operational burdens for physicists to have local access
  - Universally: fulfill the promise of opportunistic grid computing
- Technical solution is federated storage: a collection of disparate storage resources transparently accessible across a wide area via a common namespace

How it works

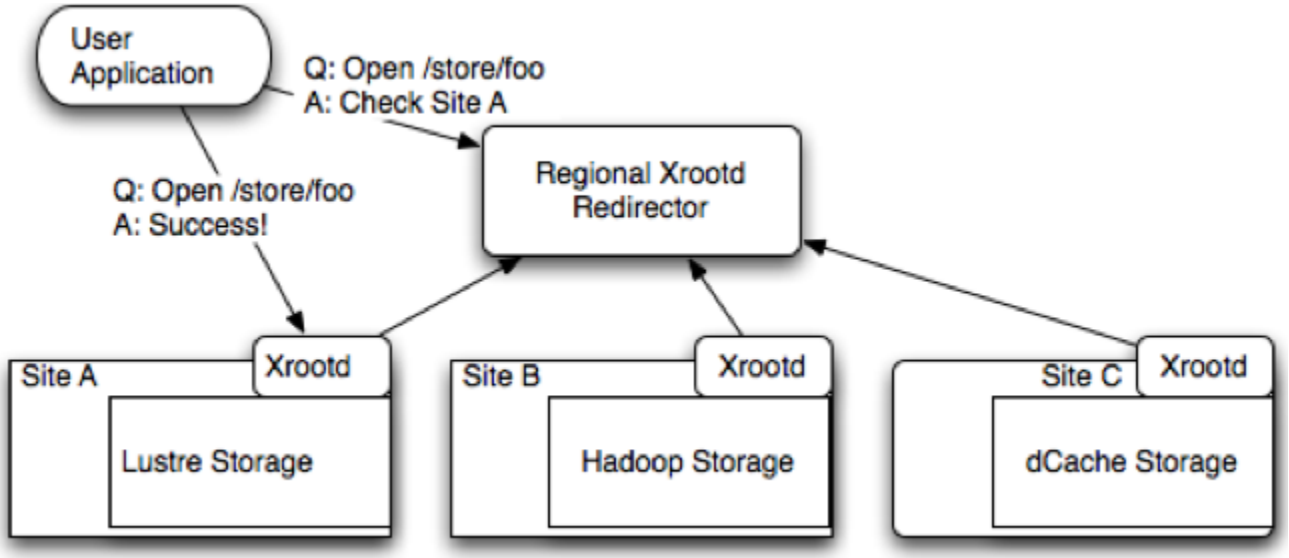

- Underlying technology is Xrootd
  - Uniform interface in front of heterogeneous storage systems
- Sites in data federation publish their data to a redirector which can then be queried by applications seeking to access data
  - If data absent in region, fall back to query for data elsewhere
- Access is authenticated

Imperial College London

Uses and how it fits in

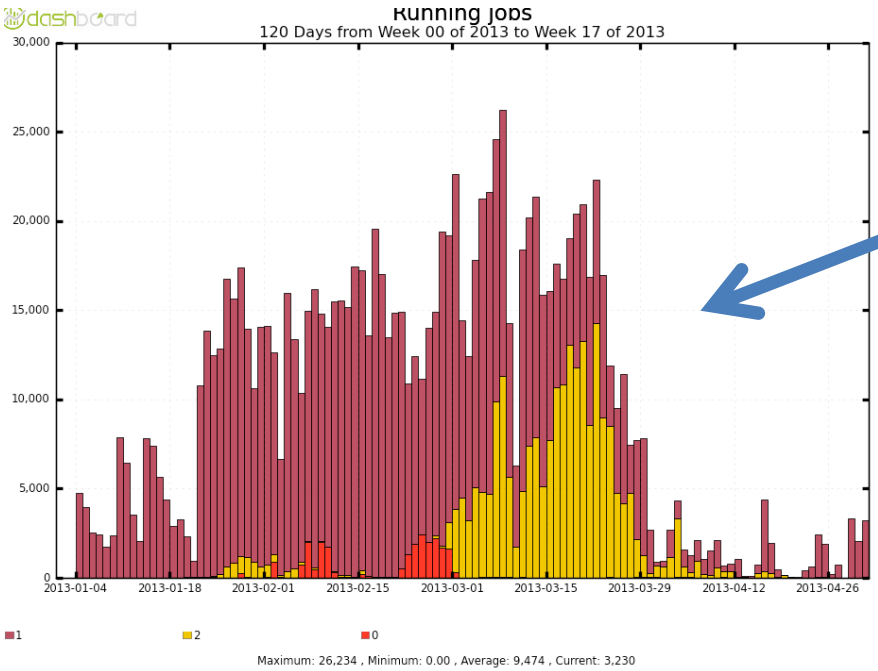

Blurring the boundary between the Tier 1 and Tier 2 sites. Example legacy re-reco. Allows efficient use of resources

Allows fewer copies of data at Tier 2 sites - saves disk

Allows diskless Tier 3 sites to access data: “At this point, I basically don’t pay attention to where the data is and just assume that jobs will find the data and run.” Kevin Lannon, Notre Dame.

Imperial College London

Current Usage

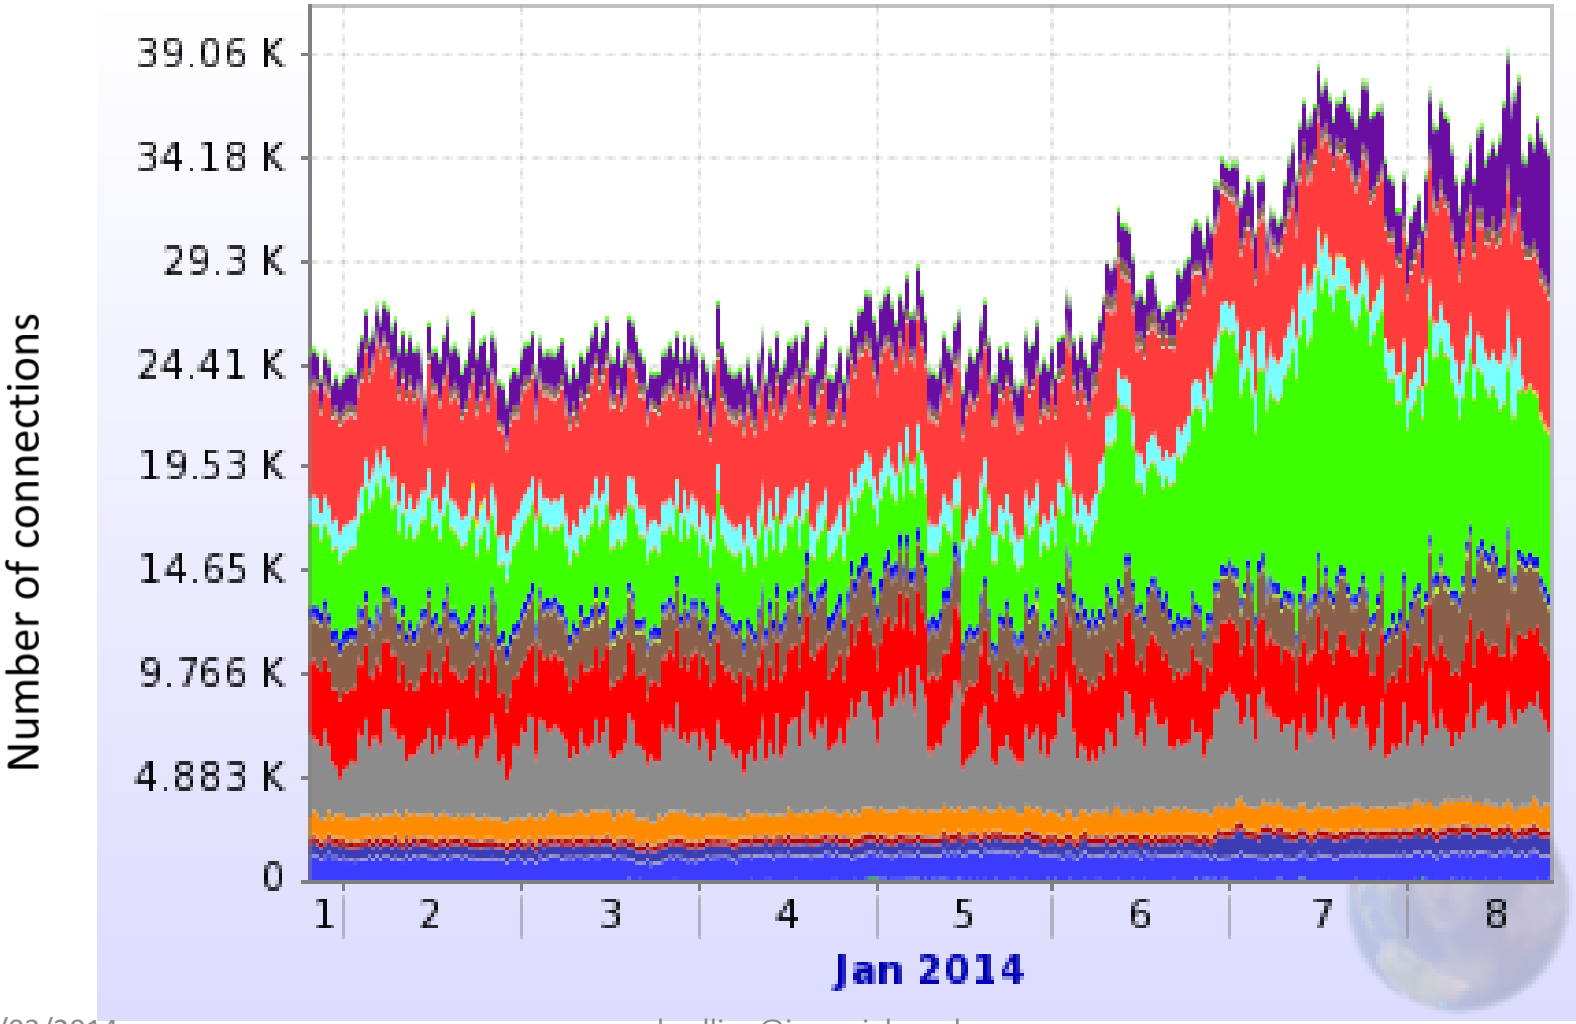

So what are going to do about it?

Further blurring between Tiers (not using AAA)

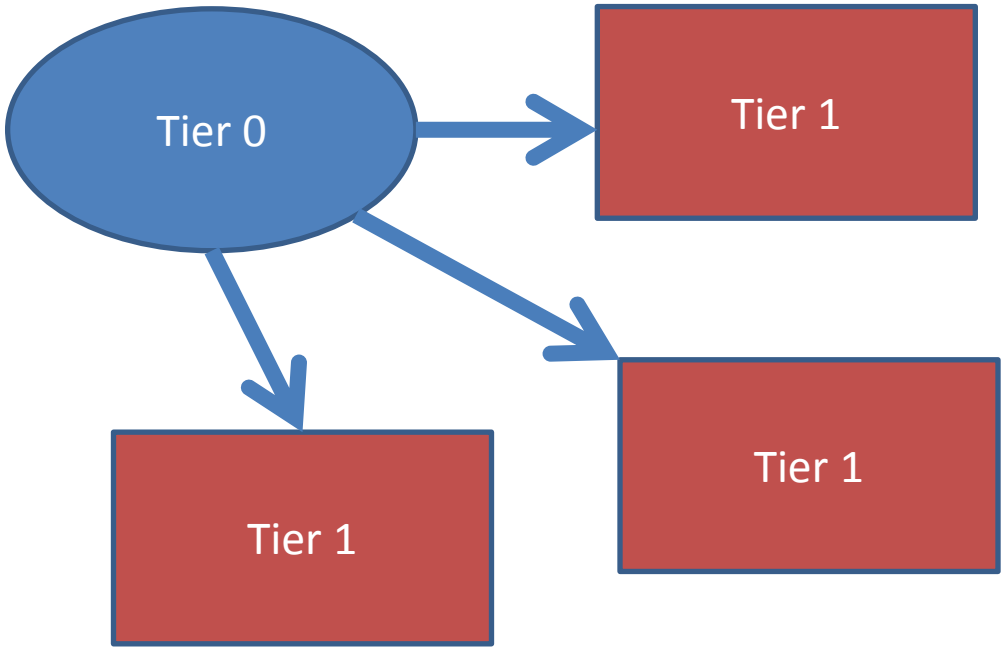

Good network links allow substantial fraction of prompt Reco to take place at Tier 1 centres

CMSooooCloud

- HLT provides CPU comparable with our total Tier 1 capacity.
- In 2013 have shown that we can use the HLT as a cloud resource during longer shutdowns (some work needs to be done here such as network upgrades but basically there)
- Using during machine maintenance is already in the planning.
- Now looking to see if we can use the inter fill period
- Even looking at times when the HLT is not fully busy.

Imperial College London

Between Stable Beams in 2012

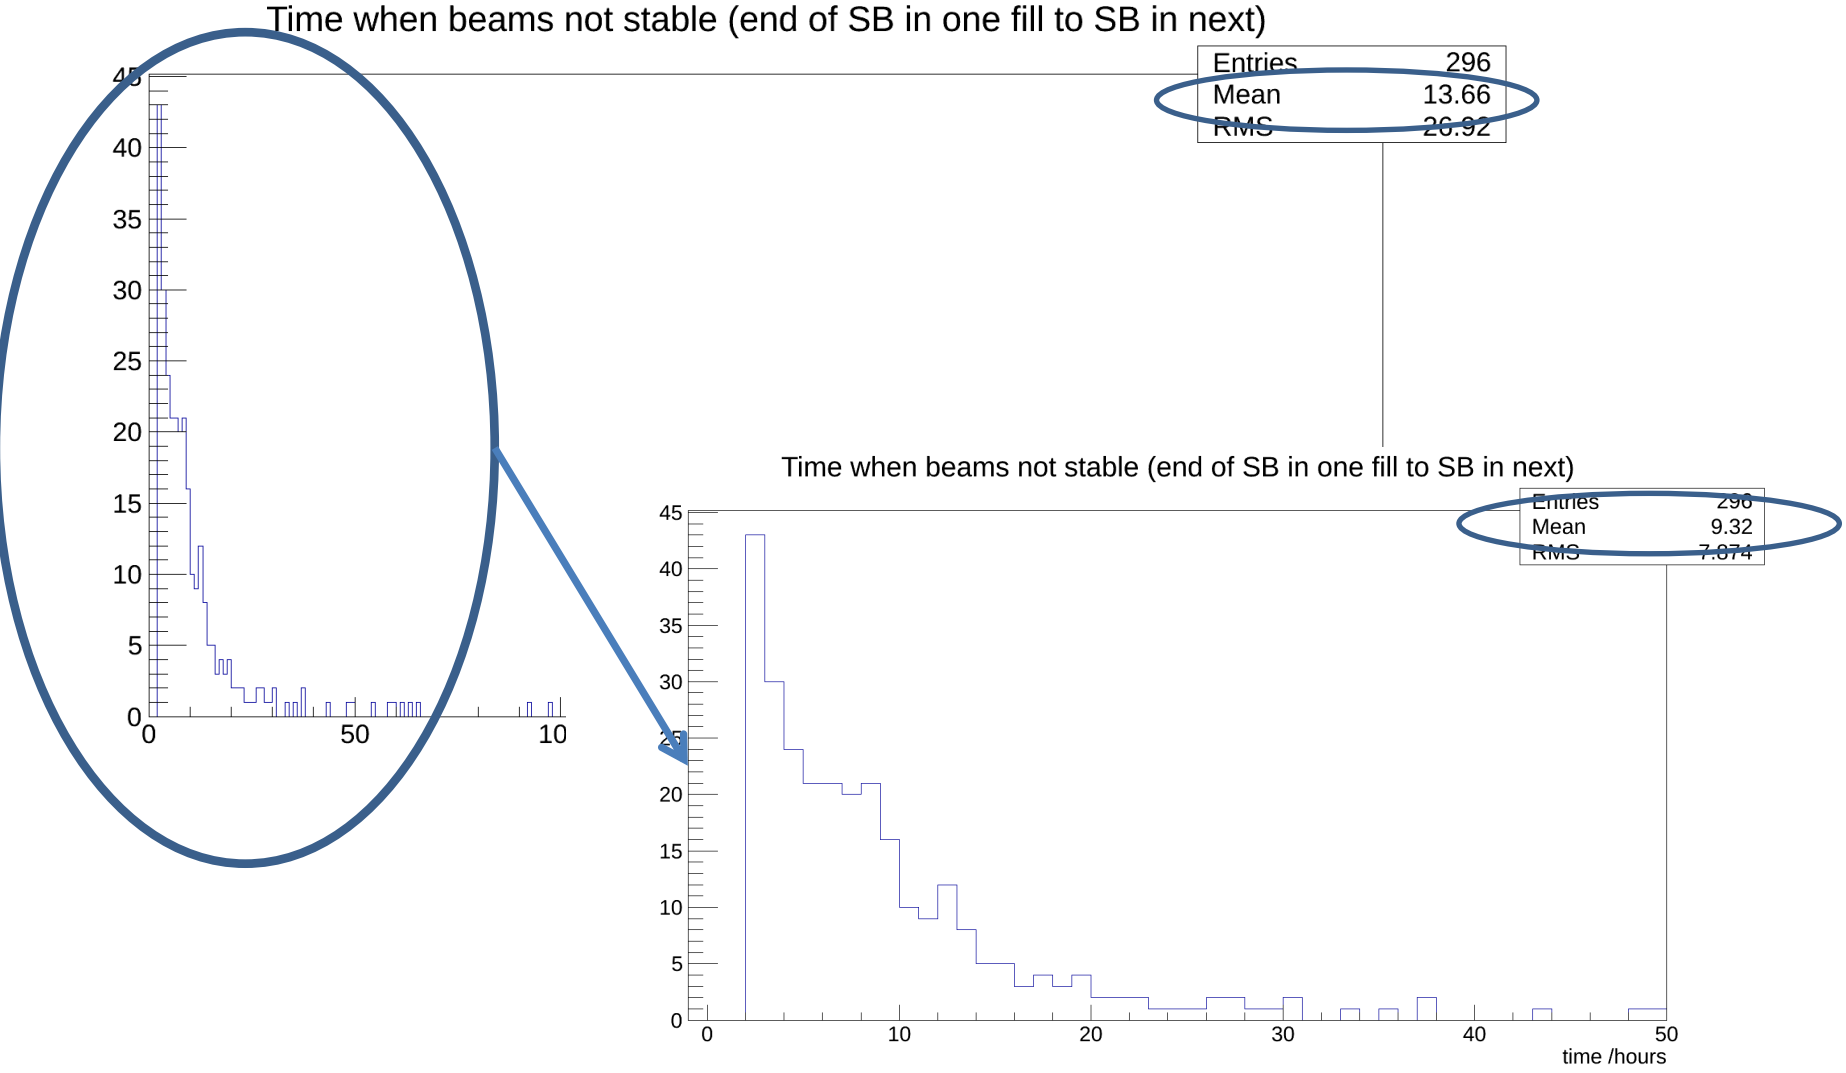

Imperial College London

Stable beams in 2012

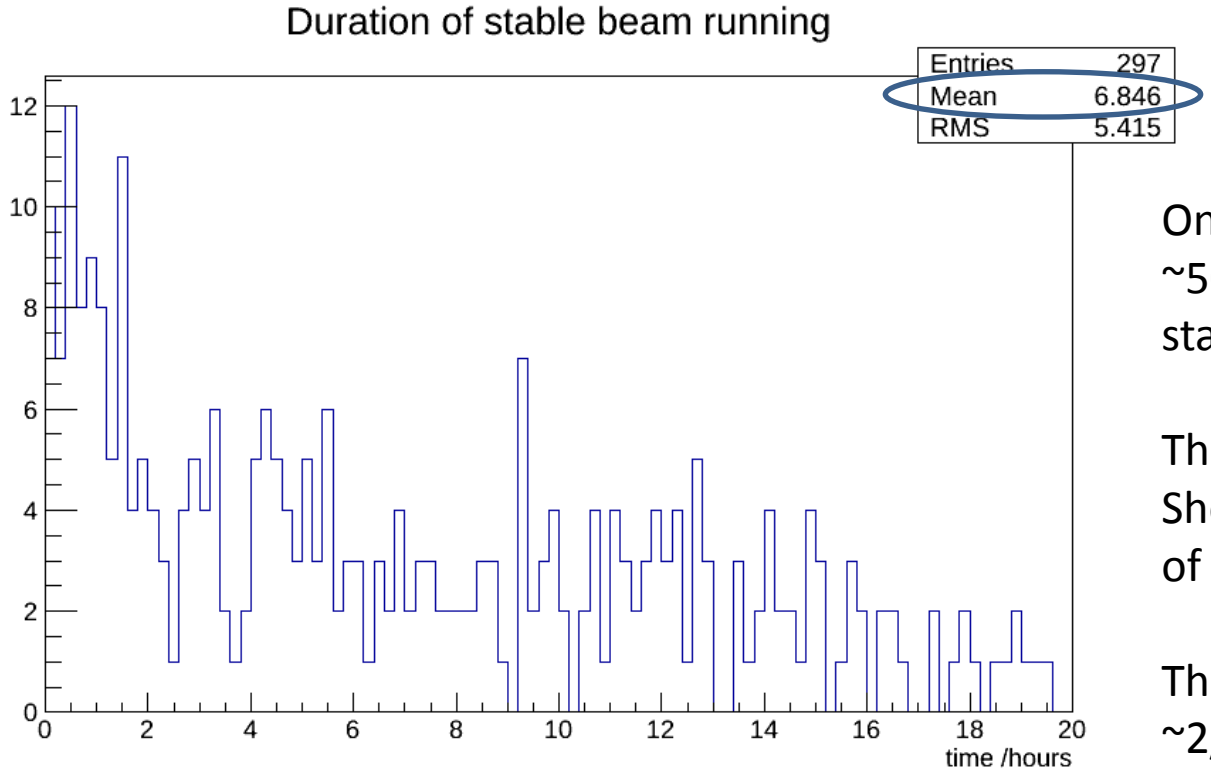

On average, stable beam is ~50% of time between stable beam.

Therefore, CMSooooCloud Should be running for >60% of 2015.

That means a resource of ~2/3 of all our T1 CPU capacity.

So, yes, it is worth trying to run between fills!

Imperial College London

Constraints on CMSooooCloud

We must never interfere with the use of HLT as an HLT

Knock on effects are:
- Online must always have ultimate control
- We can migrate "on" only when online give the green light (so we need a mechanism for this post LS1)
- We must migrate off to a state that in way interferes with the HLT.
    - o We must migrate off in a timely way - the canonical numbers are that we will get 15minutes warning.

Initial indications suggest that this will work well, now examining how to do it most efficiently.

Inter-fill running

- Tried 3 approaches
  - (default) kill and restart jobs
  - Suspend and resume
  - Checkpointing (not tried yet)

Imperial College London

Kill and restart

- Many tests performed, with much optimisation including submission rates, db cleaning etc

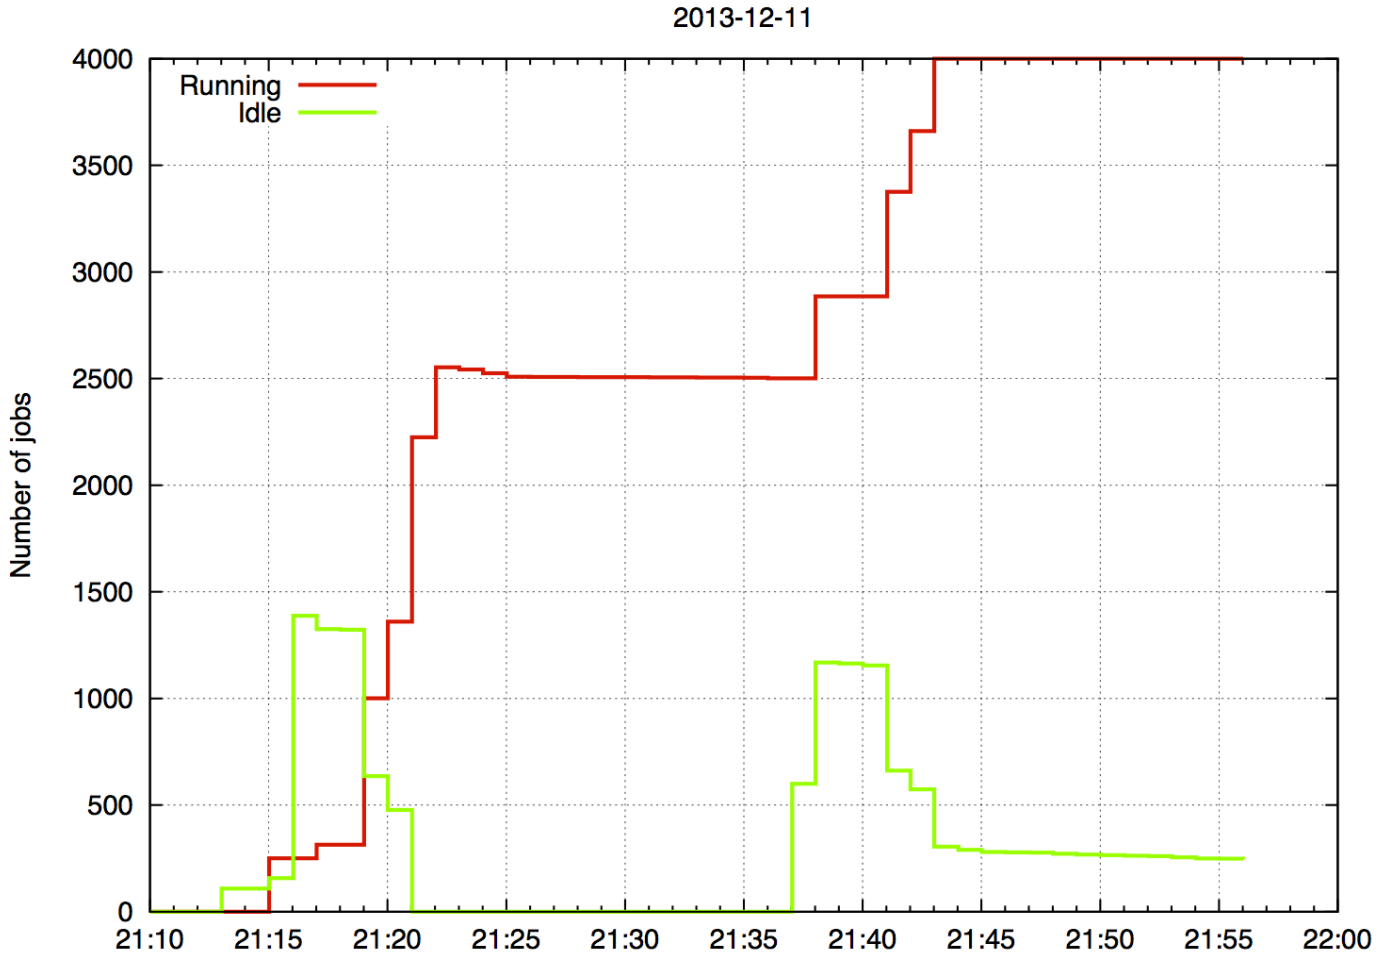

Kill and restart (pros and cons)

- (Pro) we know that it works well
- (Pro) works well with current infrastructure
- (Pro) Condor can submit the killed jobs elsewhere.
- (Con) you loose work done in the jobs that you kill.

Imperial College London

Suspend and resume

- Again tested …

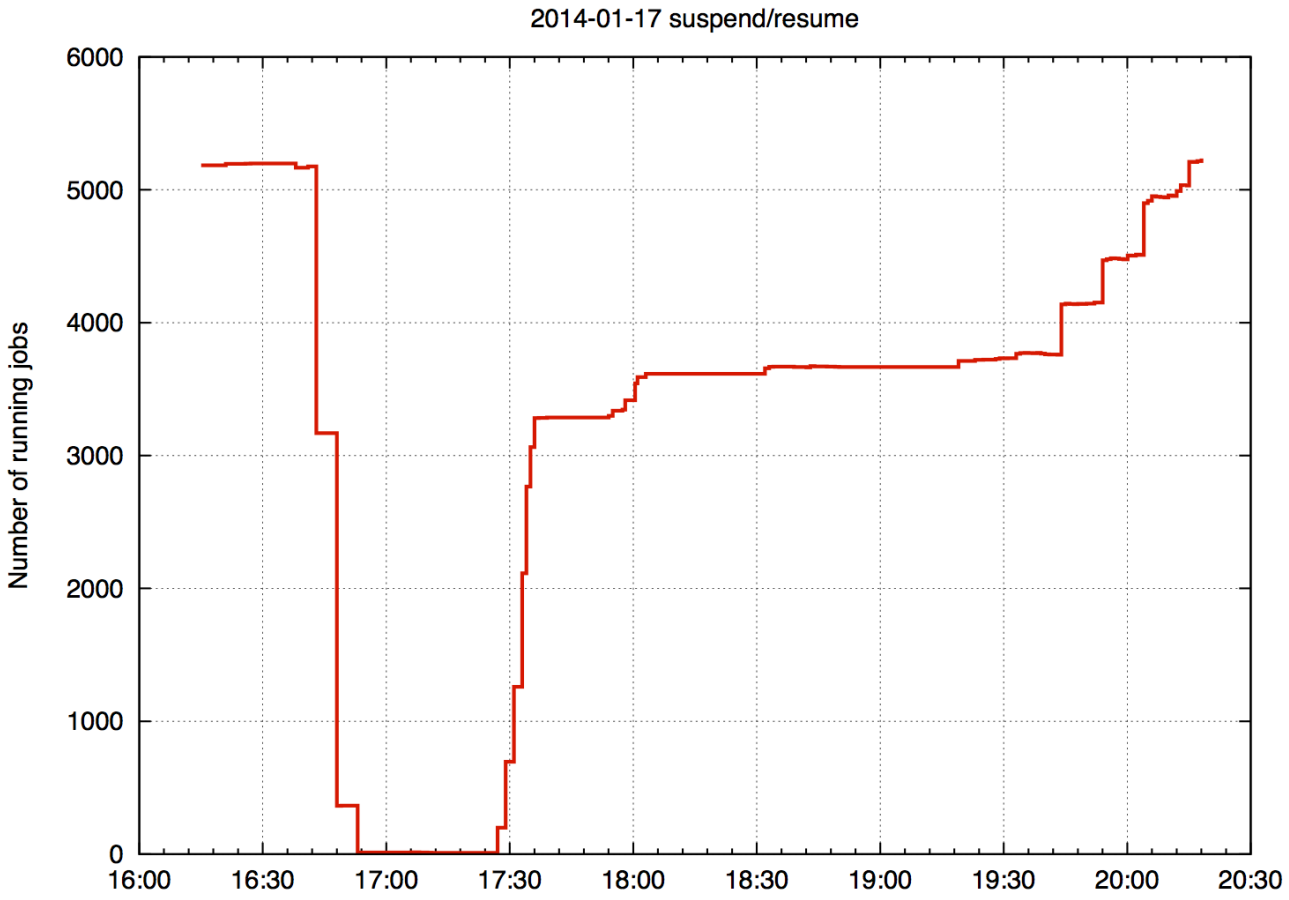

Works well and meets the time contraints.

However, not really compatible with the current production infrastructure

Checkpointing

- Lots of discussion but no action yet …

Imperial College London

Network Bandwidth

- After upgrade …

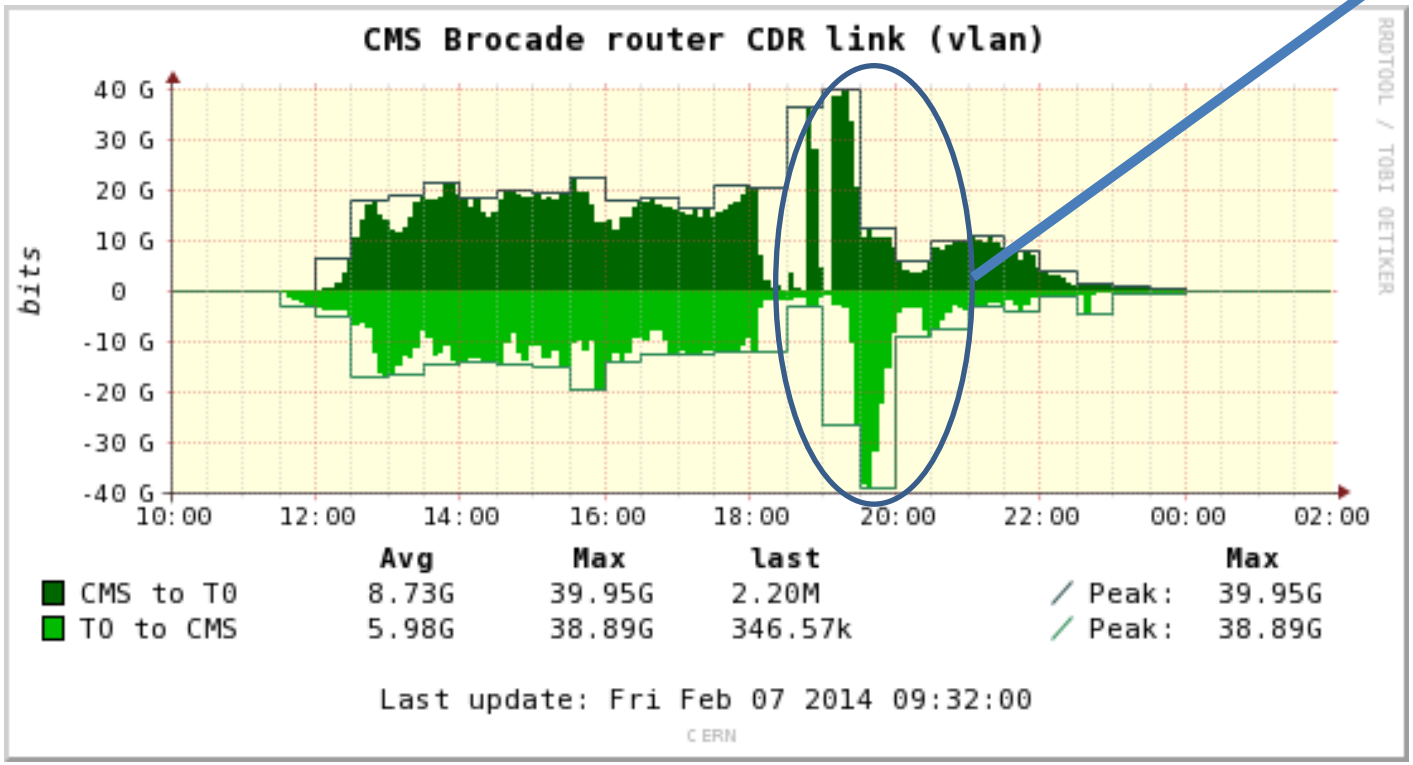

>5000 running re-reco jobs

Was bandwidth limit ed internal to cluster … now fixed.

Imperial College London

Currently … reprocessing HI data

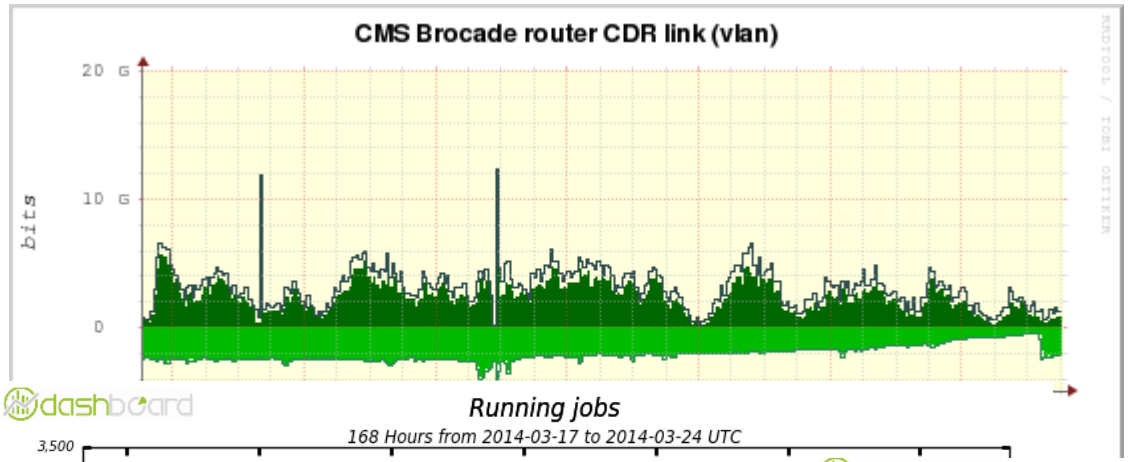

Note that all data access at over AAA

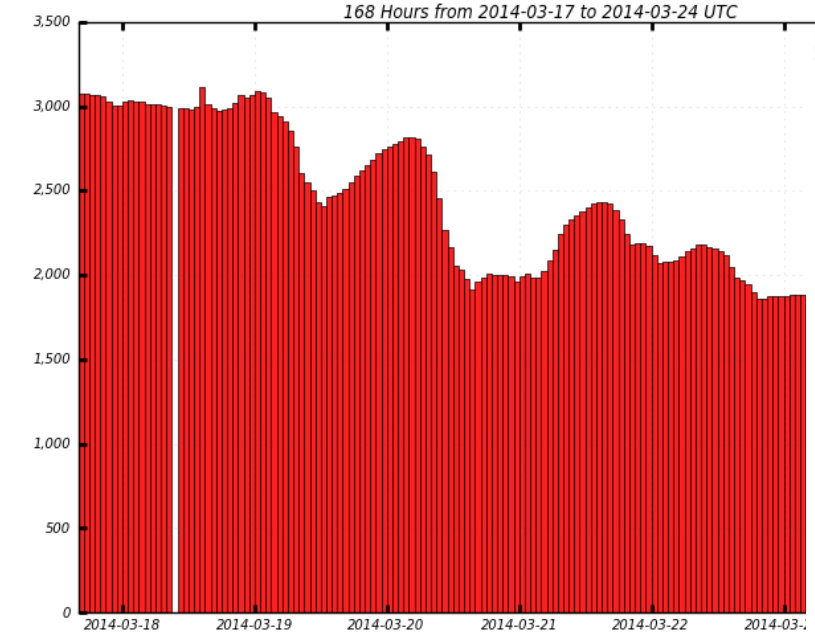

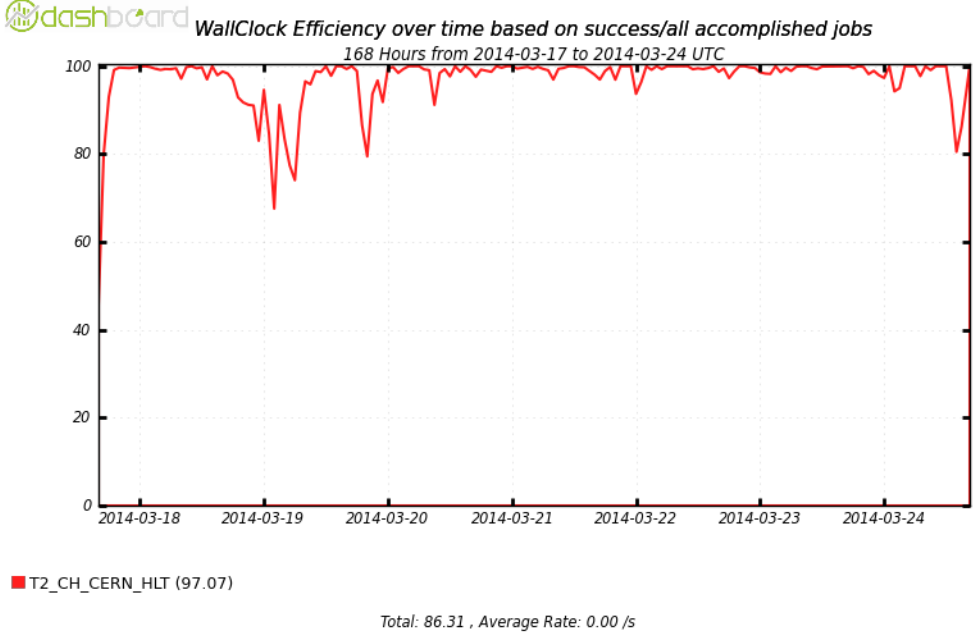

Truly opportunistic resources

- Sometimes (especially in the US) CMS is granted access to some very large resources for very limited periods of time (~few weeks)
- Often these resources are not architecturally a good match to our need – but shouldn't look a gift horse in the mouth.
- Effort is going into how to use these efficiently (using BOSCO), however these cannot be part of planning as they are unpredicatable.

The upshot

- With all this work and resulting savings and efficiency gains we retrieve a factor of ~6 (including recovery from 50->25ns)
- Still leaves a factor of ~2
- But wait, it is 2 years between the end of 2012 and 2015, so in the expected flat cash scenario Moore's Law saves us. The plan works!
- Have the research councils read the message that they are supposed to give us flat cash?