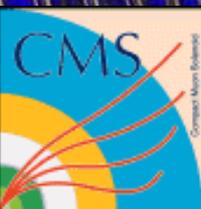


GridPP

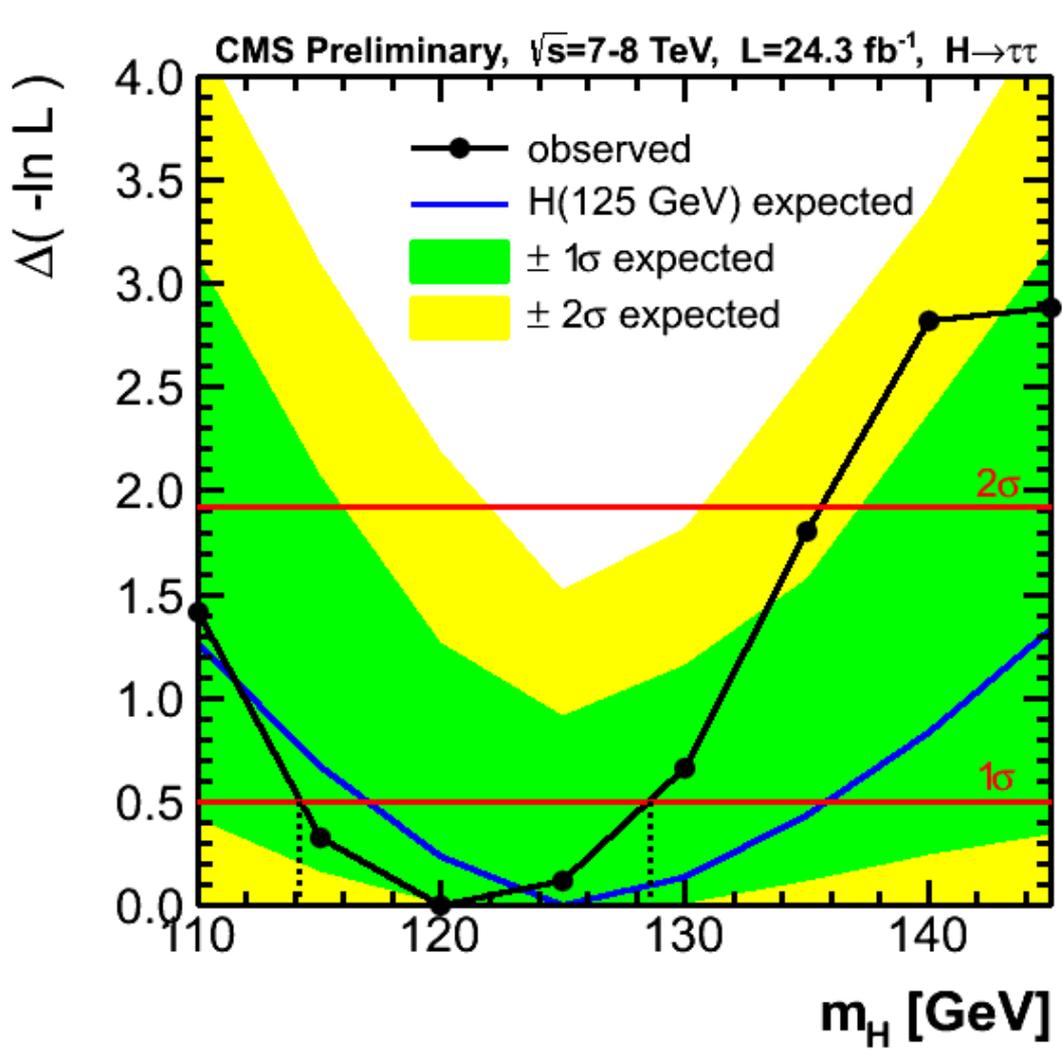
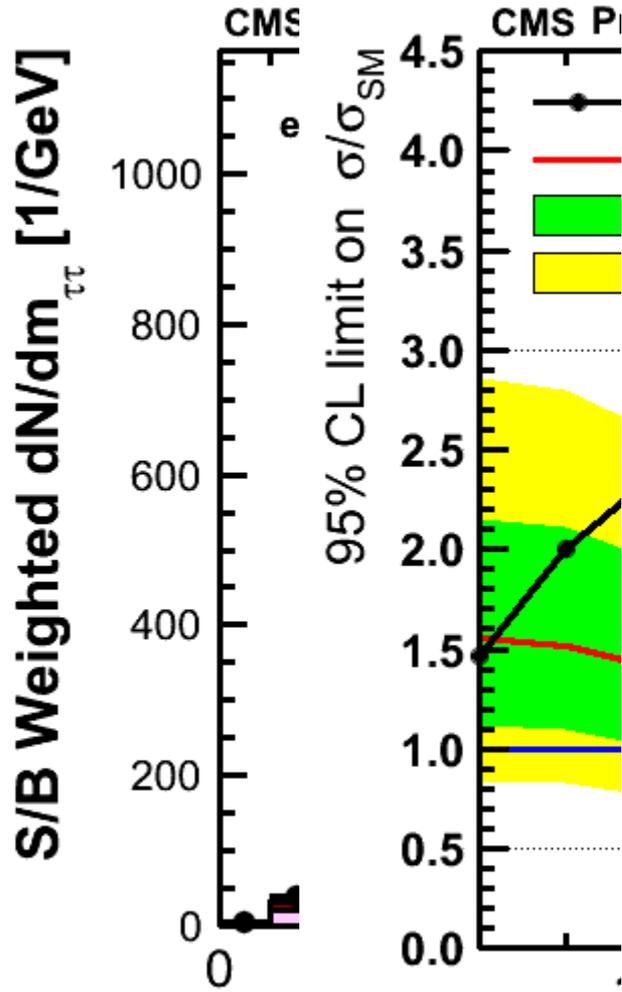
UK Computing for Particle Physics

CMS

David Colling for CMS



- Firstly thanks to all those whose talks I have plagiarised for this presentation - especially Ian, Frank, Ken, OLI, Brian, Dirk ...
- Secondly, when writing this talk I wasn't sure how much detail to go into. If it is too much detail then let me know and I will move on, if it is not enough then ask me questions.



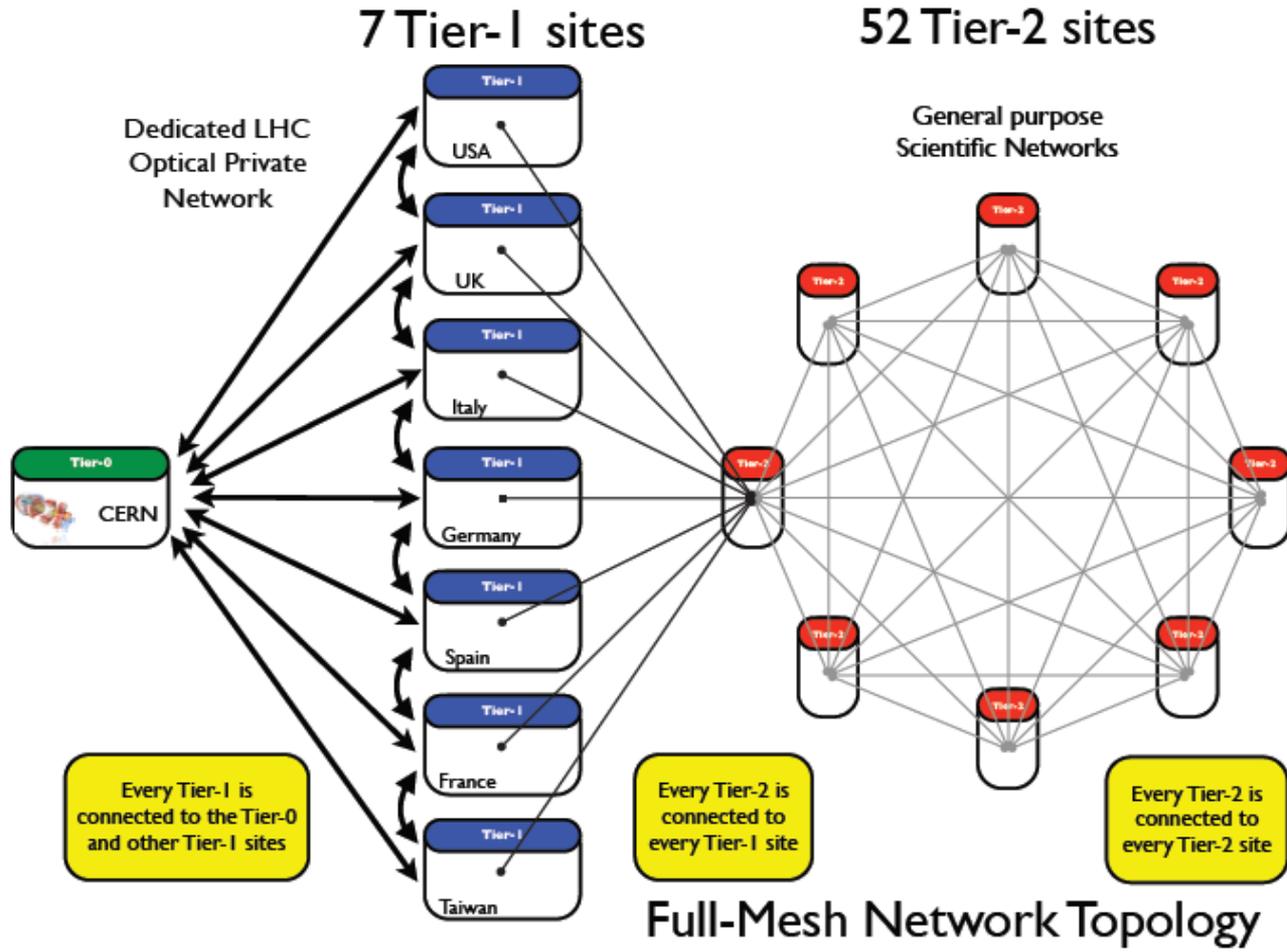
- **Pile-up is expected to grow with the increased instantaneous luminosity**
 - Roughly a factor of 2.5 increase in the reconstruction time with the best code we currently have
- **CMS currently estimates the change of the rates at similar thresholds to 2012 of 800Hz-1.2kHz of prompt reconstruction rate**
 - This is just the core samples, and would be another factor of 2.5
- **Machine will move from 50ns running to 25ns running**
 - Unexpectedly this increases the reconstruction time by roughly a factor of 2
- **Combination of all these effects is a factor of 12**

- Grids don't rule the world...
- For ~15 years x86 has been supreme ... Not for much longer
 - Rise of MIC, ARM -APU, (GP)GPUs, even FPGA accelerator cards
- Networking is much (really much) better than we had hoped in 2000
- Usage patterns not what we had first thought that they would be...

- First fix the reconstruction for 25ns (looking OK)
- Blur the Tier divides
 - Much of the prompt reco done at T1s
 - Do some reco at T2s
- Be (even) more organised
 - Data copies and placement - RECO becomes transient.
 - Number of reprocessings more like 2012 than 2010
 - Turn group analysis into a “production” activity -either group ntuple production or analysis trains
- Adapt CMSSW to run on anything that can compute
 - So able to utilise co-processors efficiently, ARM clusters etc
- Use opportunistic resources
 - HLT
 - US SuperComputers

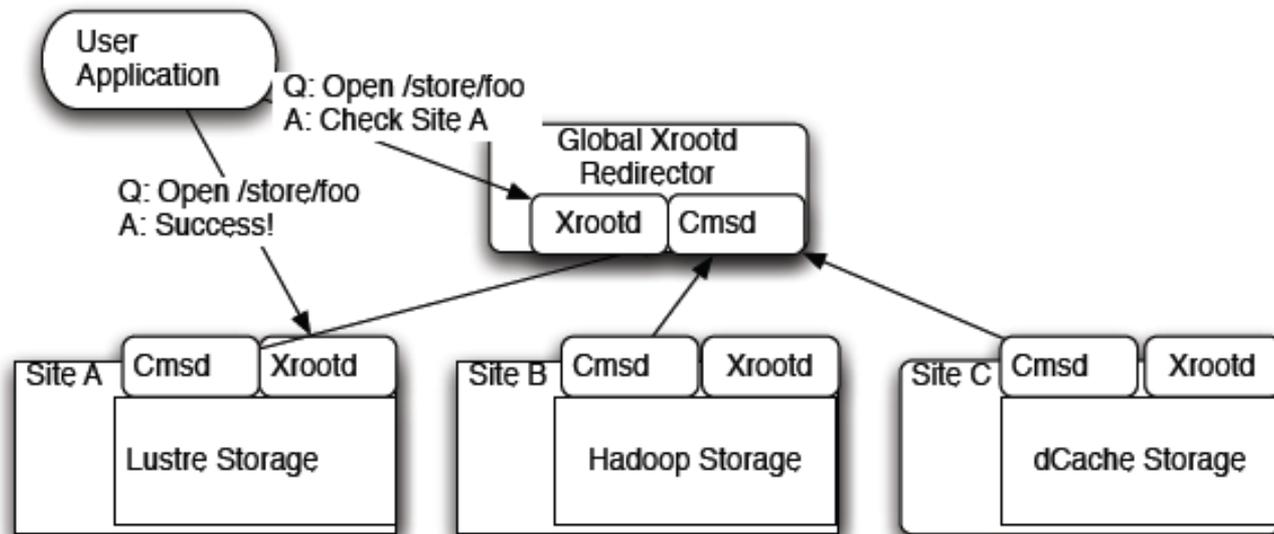
By doing all of this we hope to get about a factor of 6

... Still leaving a factor of 2 between 2012 and 2015



Network is the key!!!

- AAA
 - Any Data, Anytime, Anywhere
- Blurring Tier boundaries
- Opportunistic Resources
 - Clouds - HLT
 - BOSCO
- Multicore
 - CMSSW modifications
 - Efficiency
 - Scheduling



Global redirection system to unify all CMS data into one globally accessible namespace.

Is made possible by paying careful attention to IO layer to avoid inefficiencies due to IO related latencies.

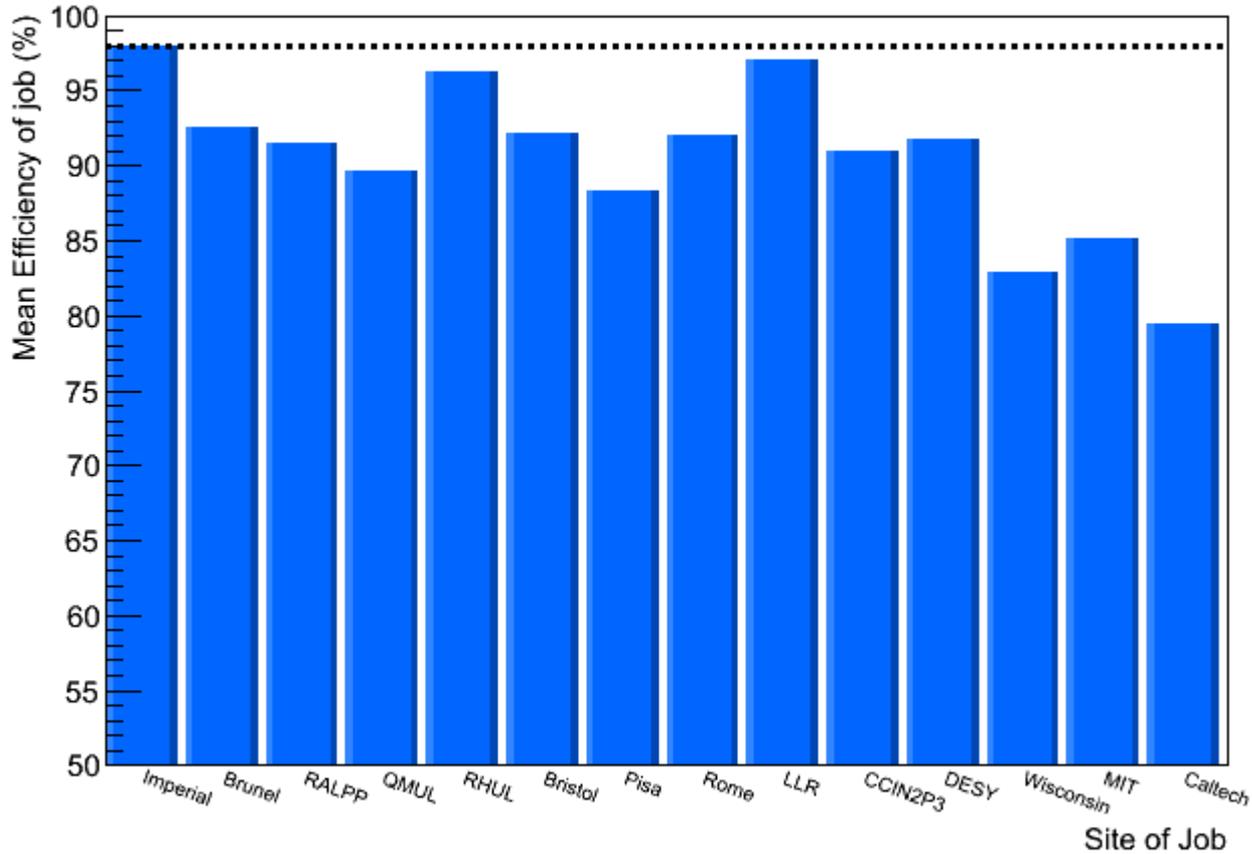
- Already widely used in production (see later)
- Used for fallback (soon) at all Tier 2 sites
 - Becomes a critical test on 1st October
- Can run user analysis at sites not hosting data
 - In the future CMS will view user black/white lists as suggestions rather than compulsory
 - Some sites much more busy than others simply because of data they host
 - More efficient to run at a “nearby” site than queue for hours
- Key feature for use of opportunistic resources
- Still being studied but will remain a key part of our plans

AAA - The Tier 3s just love it (from talk November 2012)

- ▶ We are giving physicists more control over how they work, letting them get more done, and having an impact outside CMS
- ▶ From Kevin Lannon at Notre Dame: **(to Ken Bloom)**
 - ▶ It's been a while since I promised to get back to you. This is not because I've lost interest. Quite the contrary. We've started using the AAA/xrootd mode of accessing data remotely from our T3, and I've got to say, it's really awesome. I've managed to personally run 800+ jobs all reading from various T2's (and FNAL) without any trouble. It's basically 99% reliable (much better than the luck I'd had at random grid sites). 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. So far, so good. My favorite tidbit is that head of HPC for our Center for Research Computing (CRC) tells me that when the CMS group is not running AAA jobs, the campus research traffic is about 500 Mb/s, but when we fire up our whole farm with AAA, we hit 2-3 Gb/s (and have sustained that for ~ 1 week). It's nice to be noticed. It's also a fun fact (for me anyway) that when I got here four years ago, the entire campus (research + regular network traffic) all went through a 1 Gb/s pipe. Fortunately, some of our early experiments with the T3 helped to convince them that the campus really needed 10 Gb/s. Of course, I view that as a personal challenge to saturate the new link bandwidth.

At 6.30 this morning... (Actually pretty quiet)





Data served from Imperial for same tau analysis as shown earlier

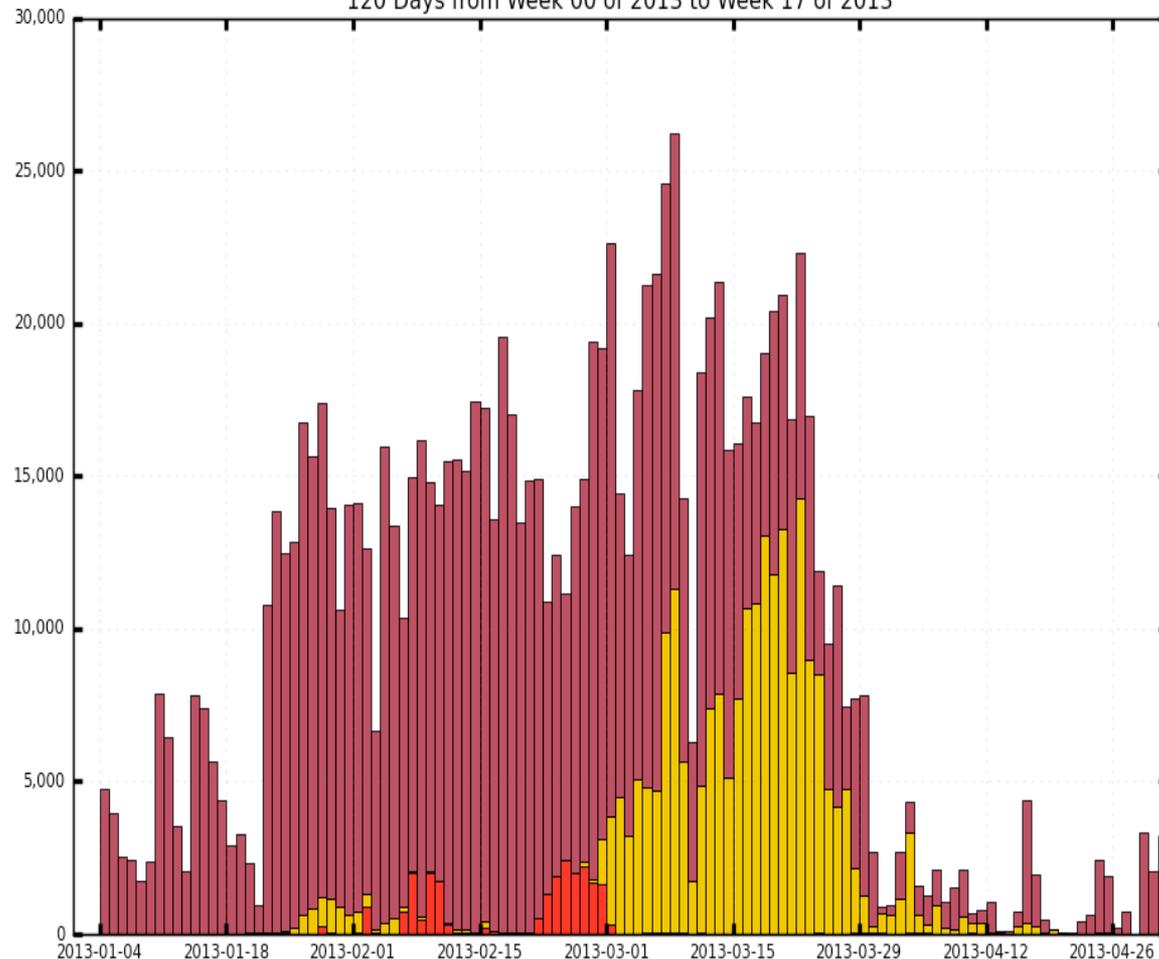


- Greater network access means that the boundary between the different tiers is blurred.
- The days of rigid structured movement/access are also blurring
 - SRMless access
 - Tasks at Tiers blurring (see next slides)
 - Make efficient use of the resources we have available



Running jobs

120 Days from Week 00 of 2013 to Week 17 of 2013



■ 1 ■ 2 ■ 0

Reprocessing
2012 data and
MC
...traditionally
only at Tier 1s

27/07/2013

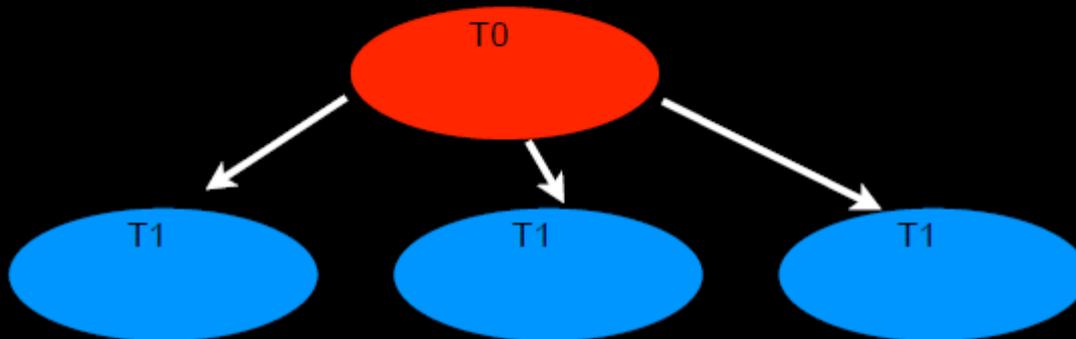




Tier-0 and Tier-1

- After Long Shutdown 1, We will likely reconstruct about have the data the first time at Tier-1s in close to real time
 - Very little unique about the functionality of the Tier-0
 - Some prompt calibration work that uses Express data, but even that could probably be exported

Already
rerecorded
data from
one Tier 1
at another.



- Just going to look at 2 (indicative) cases
 - HLT
 - UCSD supercomputer centre
- Becoming a key aspect as they offer considerable resources for limited periods

Taking effort from parts of several people so why bother?

The simple answer is that it is a big resource that we cannot afford not to use:

Node type	Number	cores/node	HS06/core	Total HS06	Disk/node (GB)
c1950	720	8	9.1	52416	72
c6100	288	12	17.3	59788.8	225
c6220	256	16	24.1	98713.6	451

- Total ~200K HS06 (of which ~150K HS06 is easily available - more than T0 and comparable with the **total** T1 cpu request)
- (Essentially) no storage available
- 2 Network paths available to CERN - 1 Gb/s Control Network, 2x10Gb/s data network
- All nodes have 2GB/core

OpenStack (Essex) installed in 2012 and initial tests with protein folding were very promising so we decided to go ahead with trying to use it for real CMS work.

The plan is to have the HLT available as a resource for (nearly) all of LS1, but then to use it as opportunistically after LS1 (in machine breaks etc, even for interventions that last more a few hours). **However, when it is need as an HLT there must be no interference from this parasitic use.** It is hoped that a cloud infrastructure will help to enable this.

The HLT was a single use cluster which meant that it didn't need the monitoring infrastructure that you would expect/need for a multipurpose

Only CMS data going from the detector to CERN IT went over the data and all other data went over the control network.

We decided to focus on reprocessing (to start with at least) and to reprocess the 2011 data.

CMS will be very short of resources in 2015 so we need to use the HLT out side of data taking. In engineering (~week long) breaks and even the in the (~12 hour) gaps between fills. Always remembering that **DATA TAKING COMES FIRST**

Even so why turn it into a cloud?

- Need to make minimal changes to the underlying set of hardware configurations when using HLT for anything else -> Virtualisation of new tasks
- Need to be able to make opportunistic use of the HLT, which means migrate on quickly and migrate off quickly (15 minutes warning) -> Virtualisation
- Complex mixture of different physical machines is not a problem as the cloud infrastructure will only instantiate VMs on resources capable of supporting them.
- Finally, potentially a good model for using other opportunistic resources as they become available.

- CMSSW served over CvmFS
- Data read from and written to EOS over xrootd
- All data read and written over 1Gb/s link
- Single frontier server installed (on cms-srv-c2c01-14)
- Submission via glideinWMS
- Images are SL5 (built with BoxGrinder)

Found many, often minor but annoying, problems.

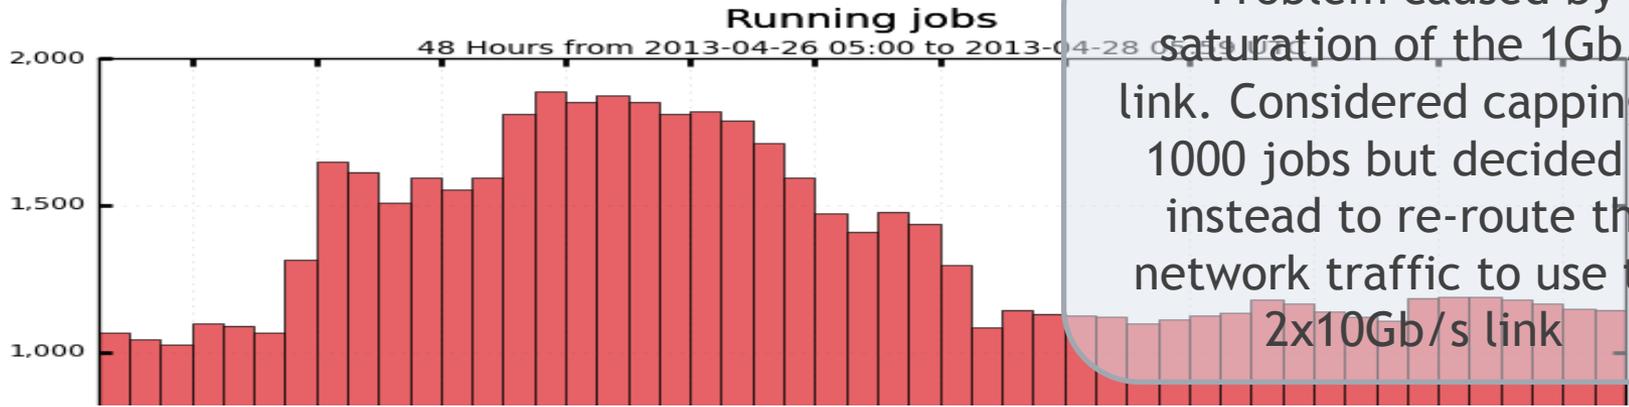
These include:

- Permissions problems with xrootd and EOS
- VMs dying because access to CvmFS was not available fast enough
- OpenStack EC2 not Amazon EC2 causing many minor problems all of which required modifications to the glideinWMS.
- Behaviour in clouds is different from behaviour in Grids so glideinWMS needed to learn how to handle the situations differently
- OpenStack controller can be “rather fragile” when asked to do things at scale so glideWMS learnt to treat it gently.
- glideinWMS losing track of jobs (often through fragility of OpenStack) and jobs ending up in “shutoff” state
- ...

Gradually been working our way through these with the set up becoming functional and then more robust as we go...

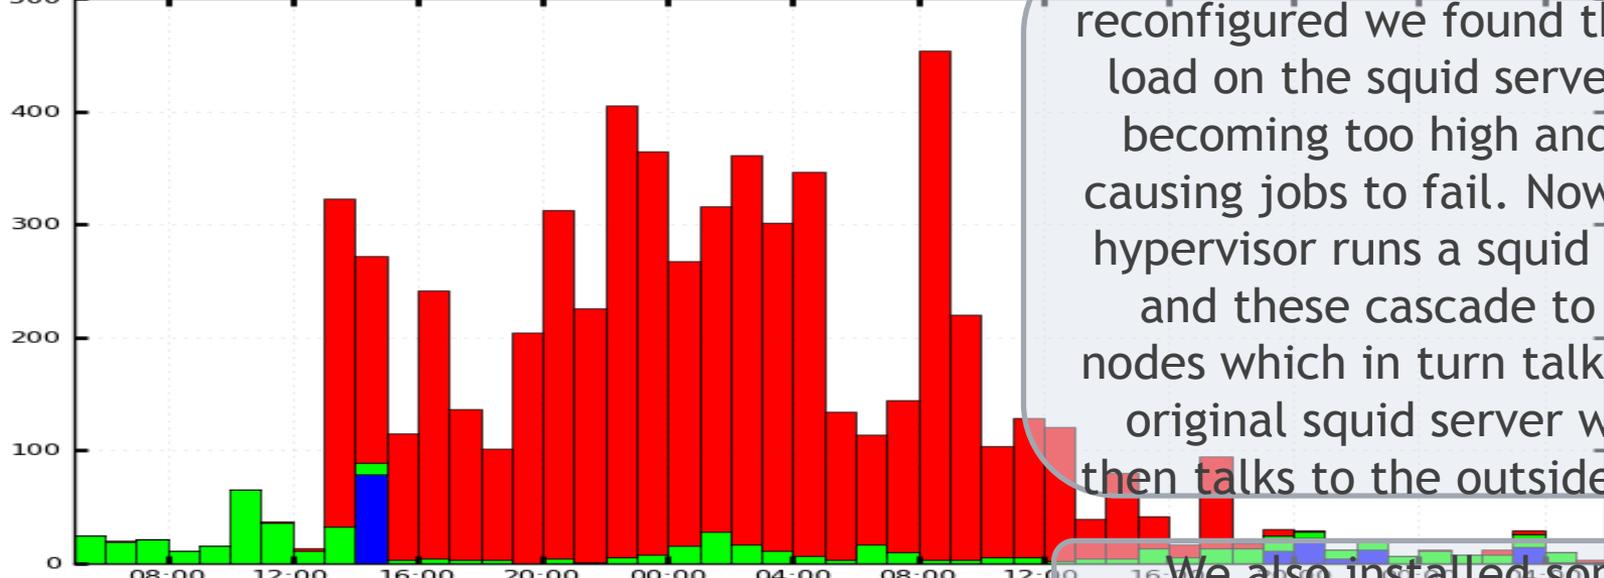


Initial results - basic limitations



Problem caused by saturation of the 1Gb/s link. Considered capping at 1000 jobs but decided to instead to re-route the network traffic to use the 2x10Gb/s link

Application Status of Terminated Jobs (Time Spent in Error)



Once the network had been reconfigured we found that the load on the squid server was becoming too high and was causing jobs to fail. Now, each hypervisor runs a squid server and these cascade to two nodes which in turn talk to the original squid server which then talks to the outside world

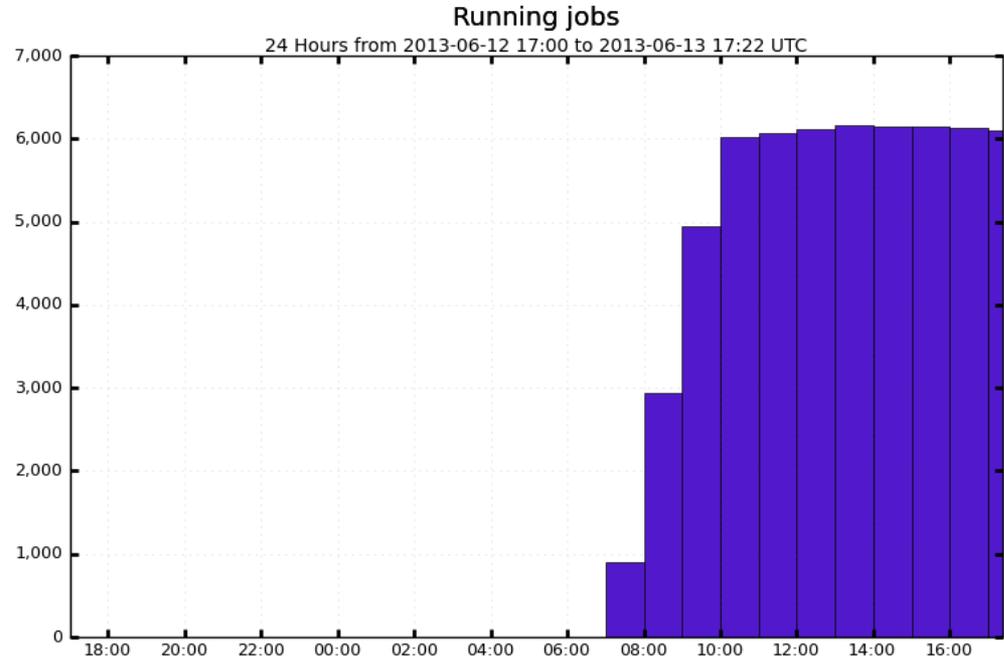
■ Number of Failed Jobs ■ Number of Successful Jobs ■ Number of Unknown-Status Jobs

Maximum: 454.00 , Minimum: 3.00 , Average: 130.43 , Current: 3.00

We also installed some monitoring

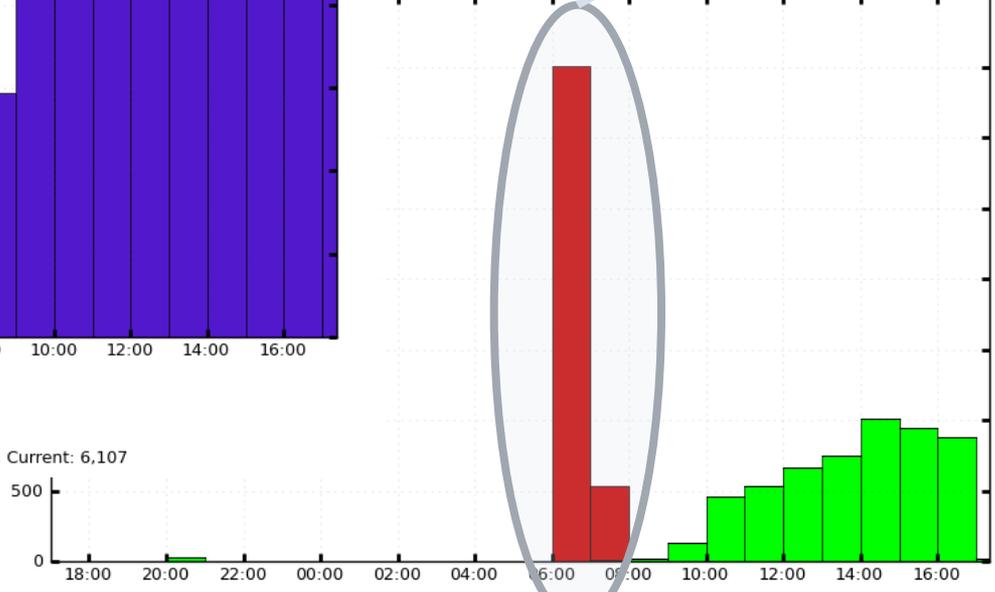


As of 13th June we are running with a new pre-release of condor and ...



Workflows submitted with wrong requirements, were cancelled

Failed and Cancelled Jobs (Time Stacked Bar Graph)
2013-06-12 17:00 to 2013-06-13 17:20 UTC

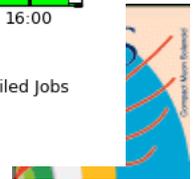


■ T2_CH_CERN_HLT

Maximum: 6,164 , Minimum: 0.00 , Average: 2,308 , Current: 6,107

■ Number of GRID-Failed Jobs ■ Number of Successful Jobs ■ Number of Application-Failed Jobs
■ Number of Unknown-Status Jobs

Maximum: 3,503 , Minimum: 0.00 , Average: 378.96 , Current: 15.00



Number of Successful and Failed Jobs (Time Stacked Bar Graph)

24 Hours from 2013-06-13 10:00 to 2013-06-14 10:03 UTC

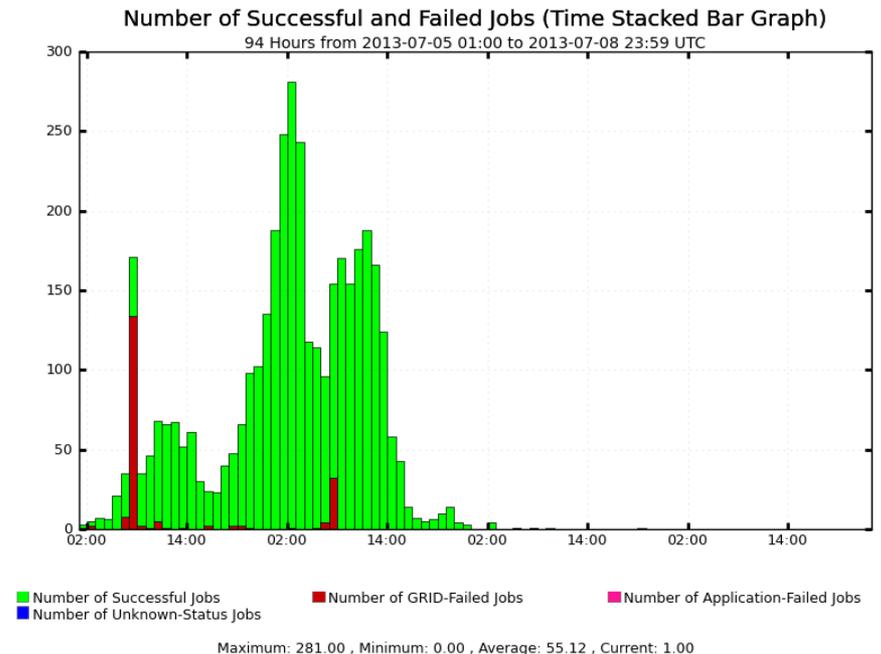
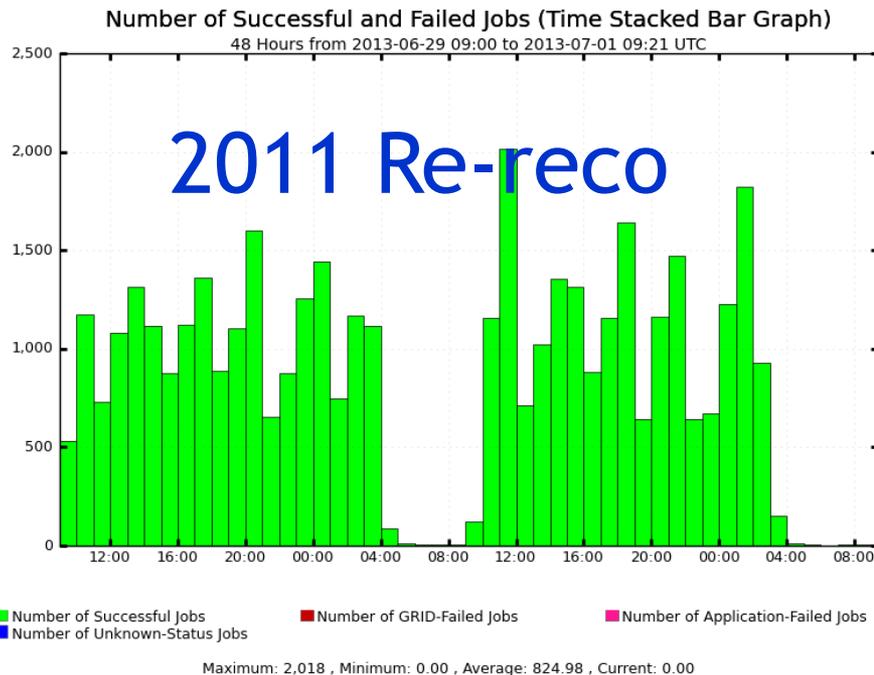


Ran a steady 6000 jobs overnight.

Some problem at ~7am (being investigated) but did recover, still 97% successful!

Maximum: 1,347 , Minimum: 96.00 , Average: 769.58 , Current: 96.00

- Being used as a production resource

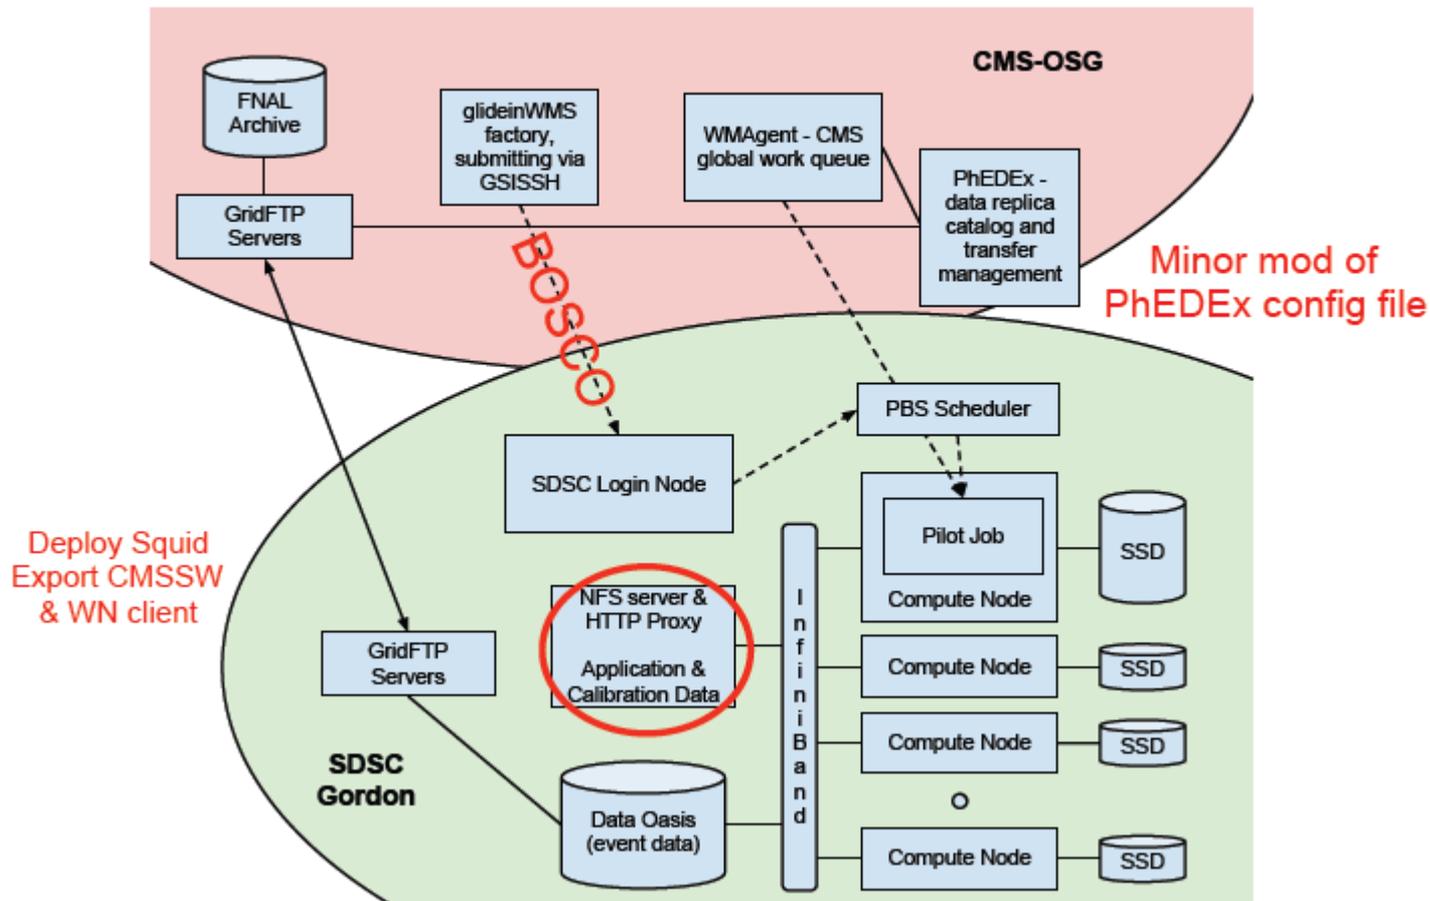


Re-digi MC samples

- Usable production service - however
 - Need to upgrade network, still measuring by how much
 - Will upgrade to (easily usable) 200KHS06 during LS1
 - Starting work on migration strategies - got to get off in 15minutes
 - Clearly a key part of CMS future resources

Example of the GORDON supercomputer at UCSD

Items in **red** were deployed/modified to incorporate Gordon



- CMS will be short of computational resources in 2015
 - Real chance that data taking will be limited by what we can afford to spend on computing
- CMS will use whatever resources we can get our hands on
- Even then need to be very efficient and clever
 - The network is the key to be able to do this
- Need to modify software to work on new architectures (not covered in this talk)

GridPP people are playing key roles.

(especially Andrew L., Adam H. - being joined by Daniela B and Alex R.)