CERN-RRB-2011-106

Project Status Report

Ian Bird

Computing Resource Review Board

18th October, 2011

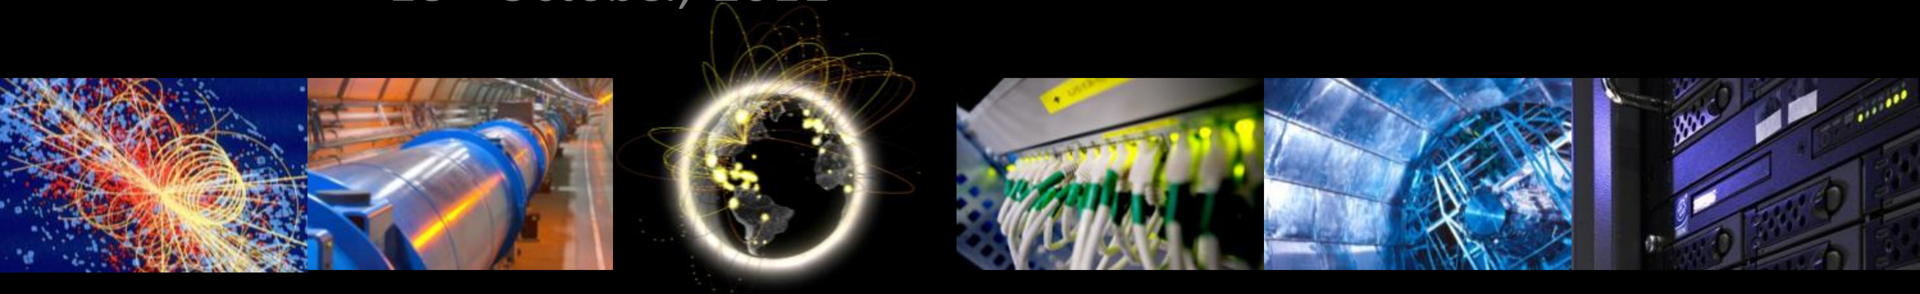

WLCG Worldwide LHC Computing Grid

CERN

Data in 2011

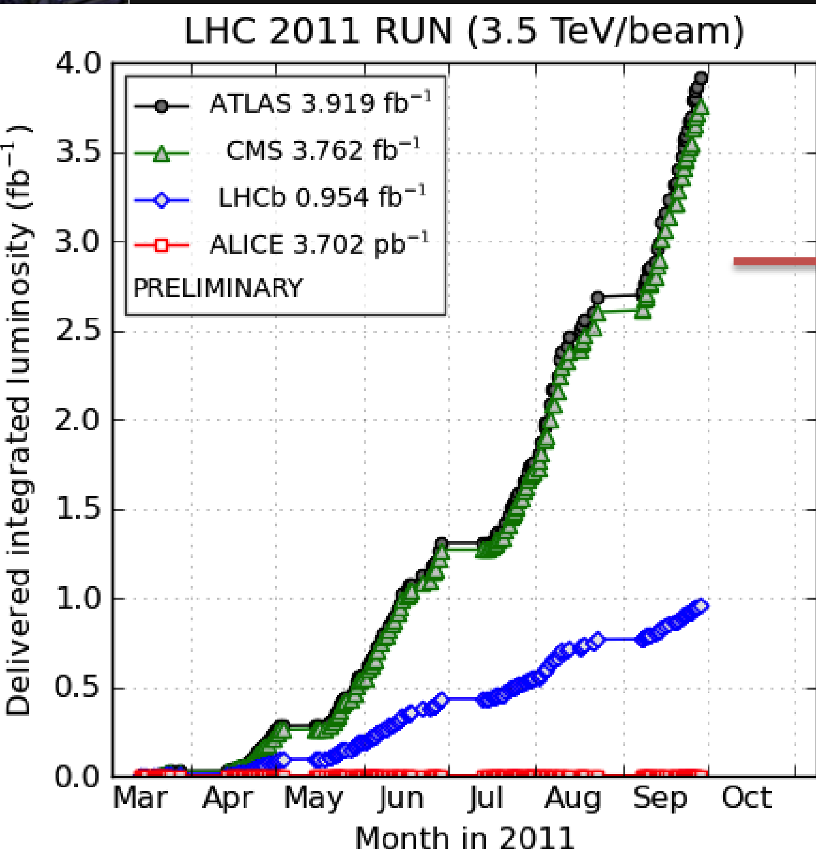

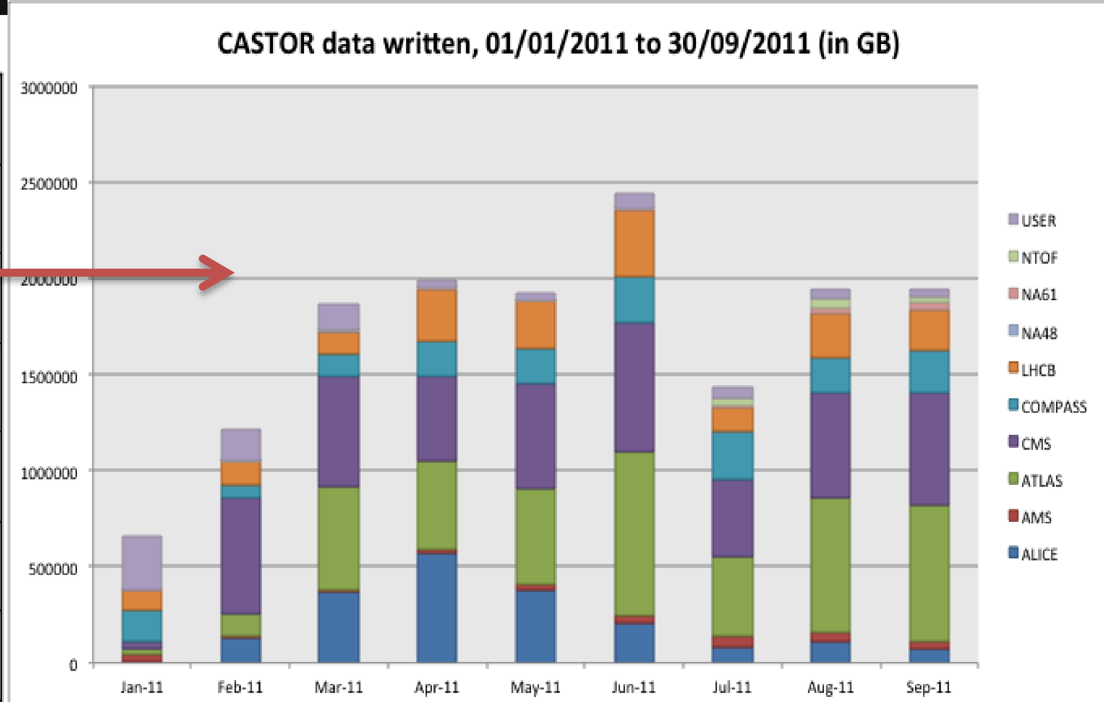

Total data written Jan-Sep: ~15 PB

… heading for 20 PB?

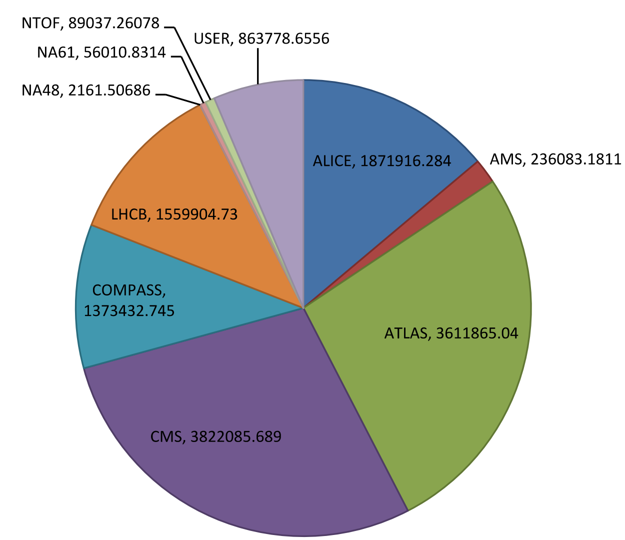

CASTOR data written, 01/01/2011 to 31/08/2011 (in GB)

Data in the Tier 0

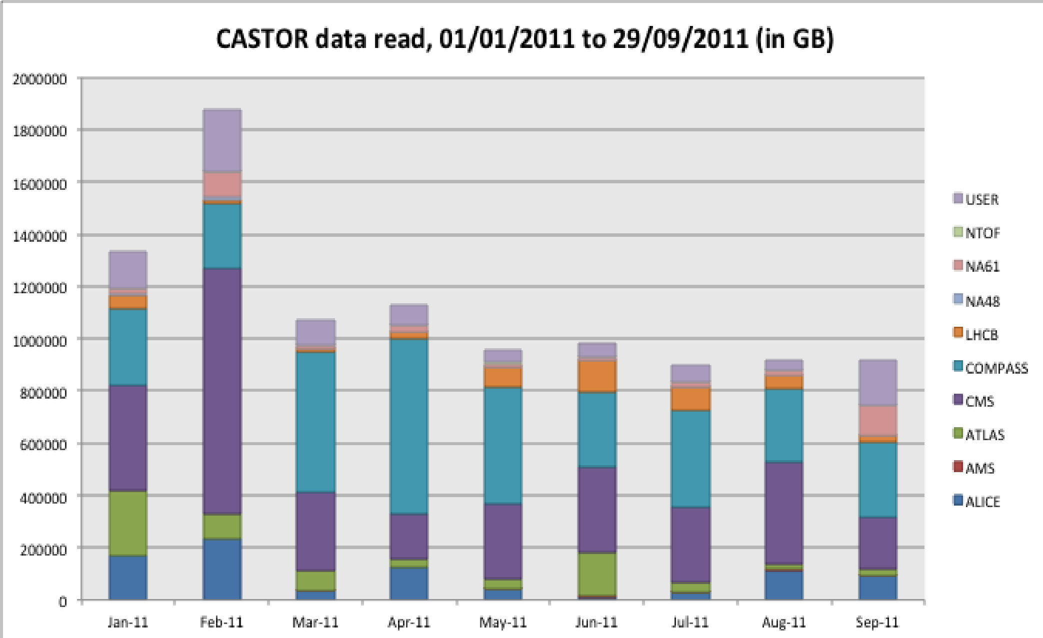

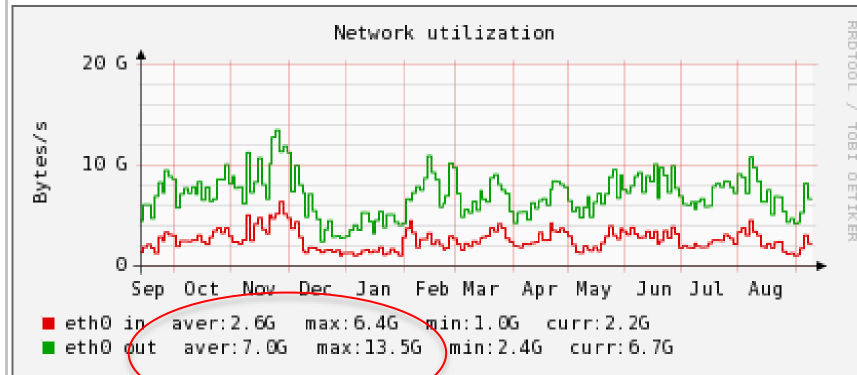

Data traffic in Tier 0 and to grid similar to 2010 values:

Up to 4 GB/s from DAQs to tape

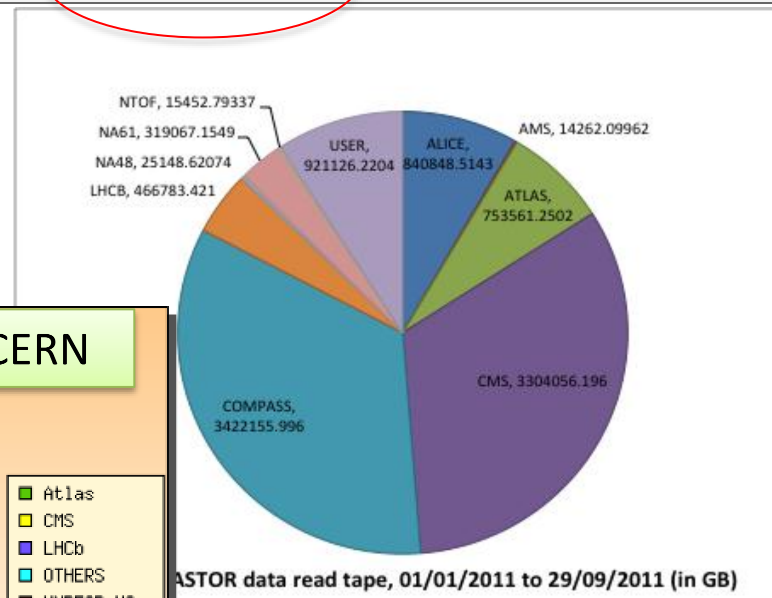

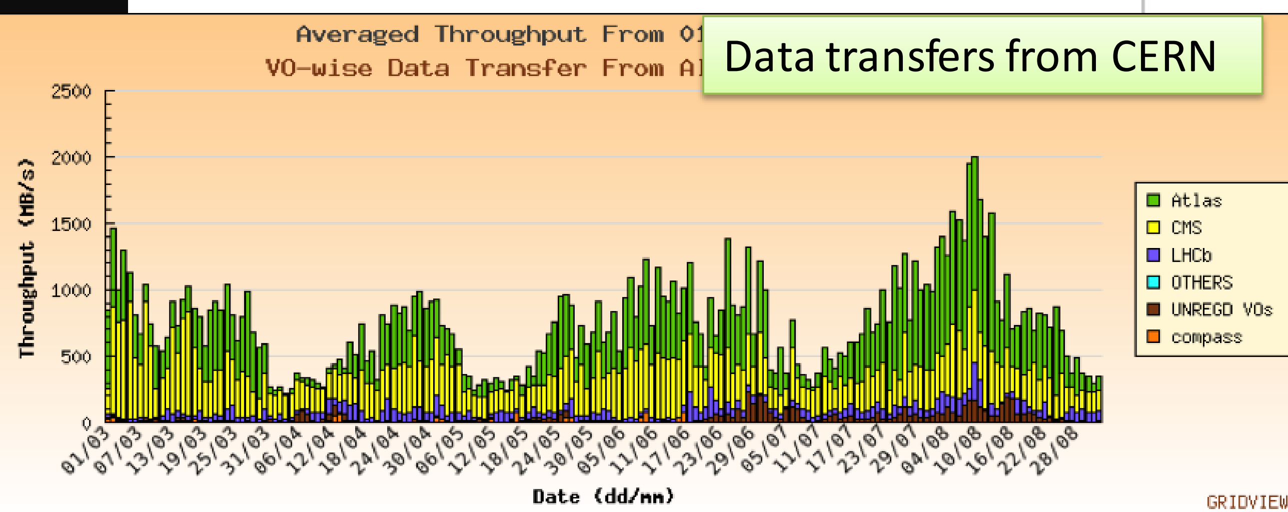

Data transfers from CERN

Tape data read; now far more controlled

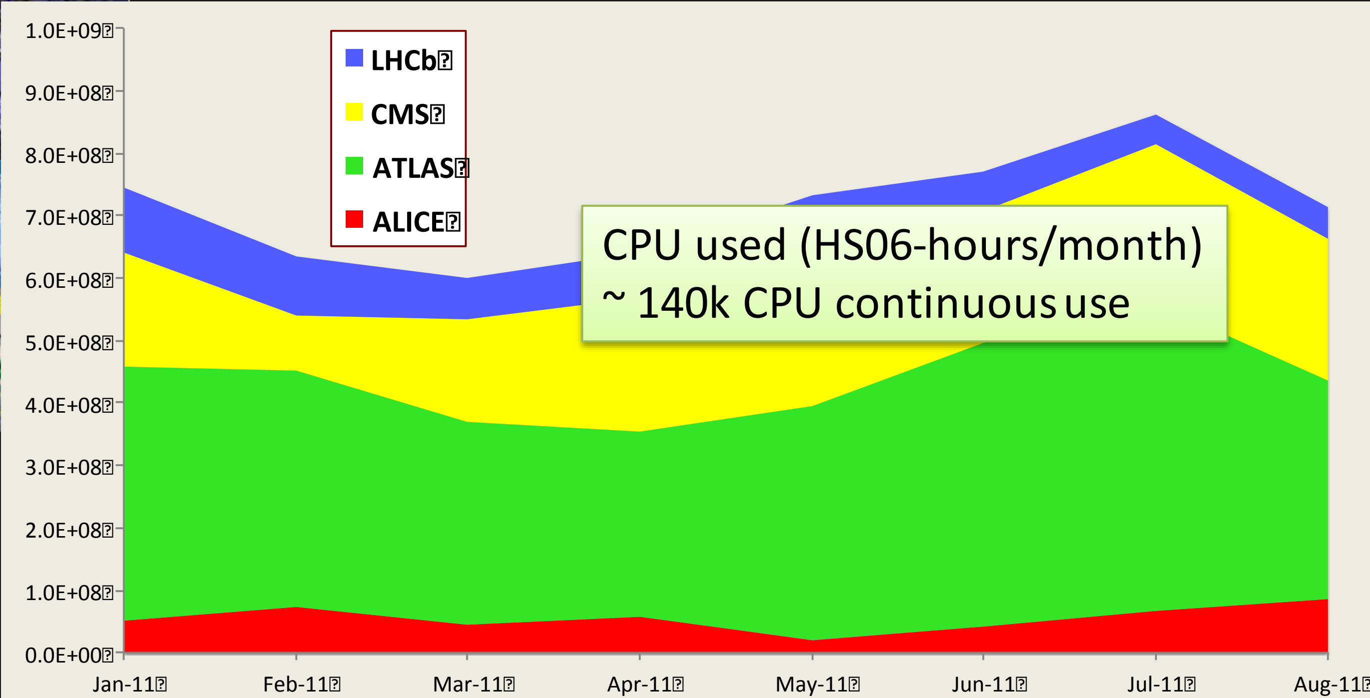

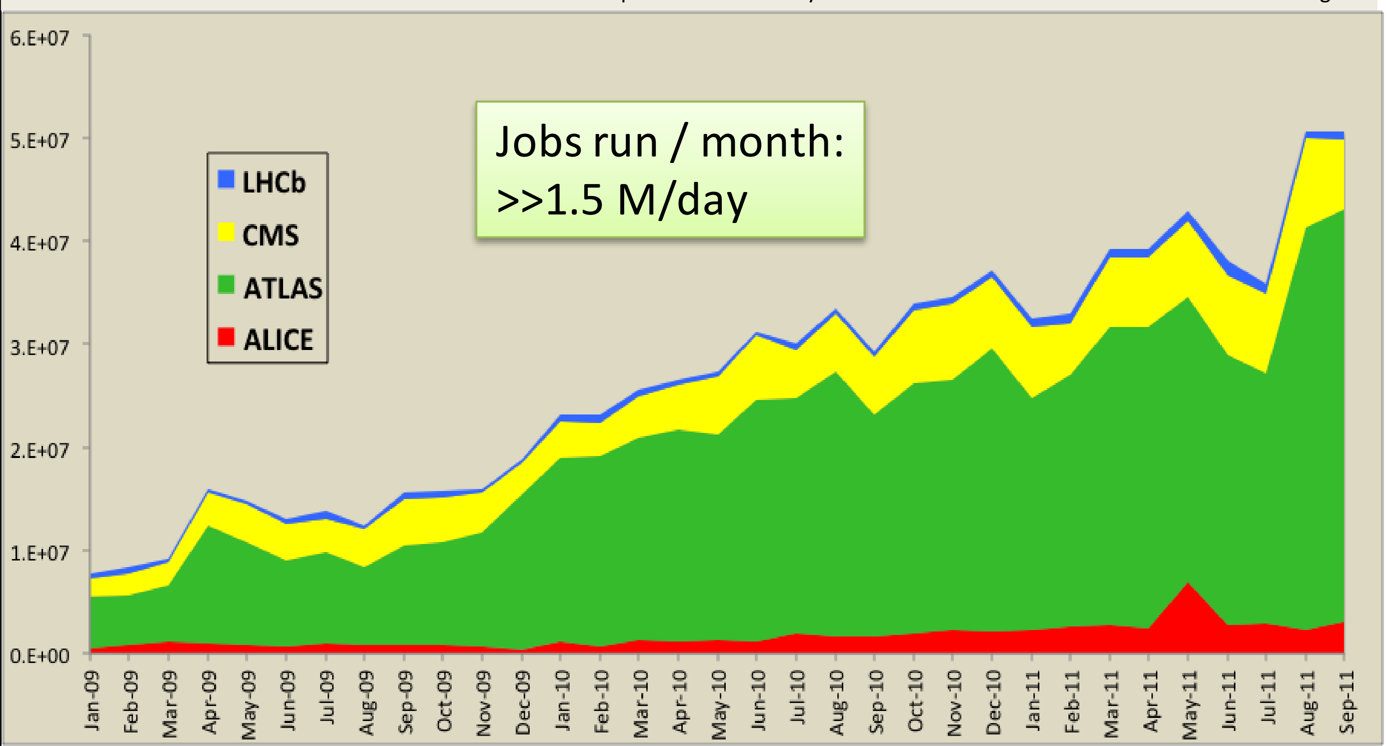

WLCG … business as usual:

- Large scale (>> TDR) both workloads and data transfers
- Resources become fully used

Summary of usage

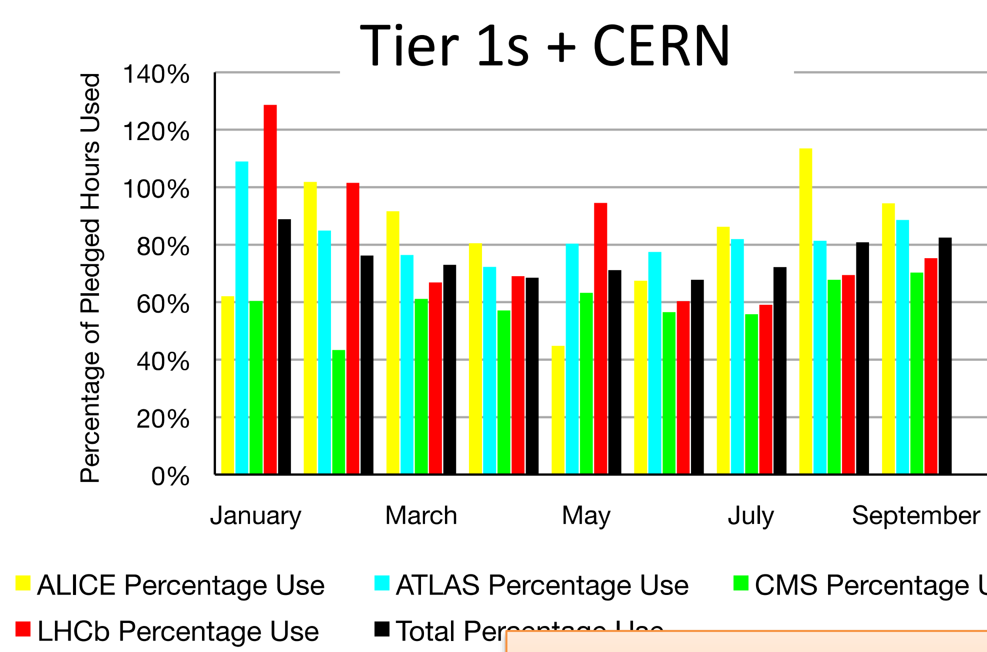

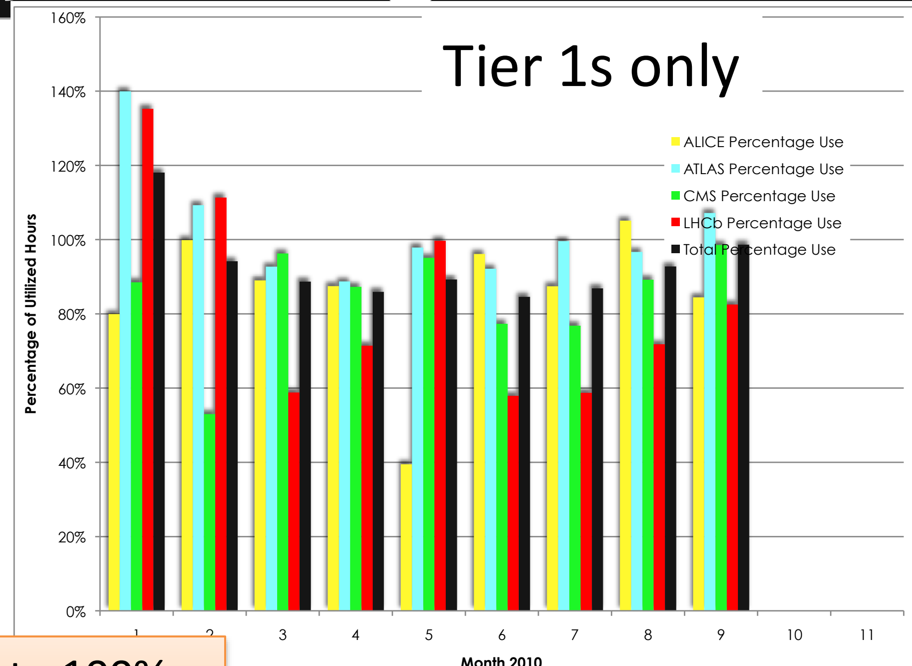

Tier 1s and Tier 2s close to 100% occupation; with generally good efficiency

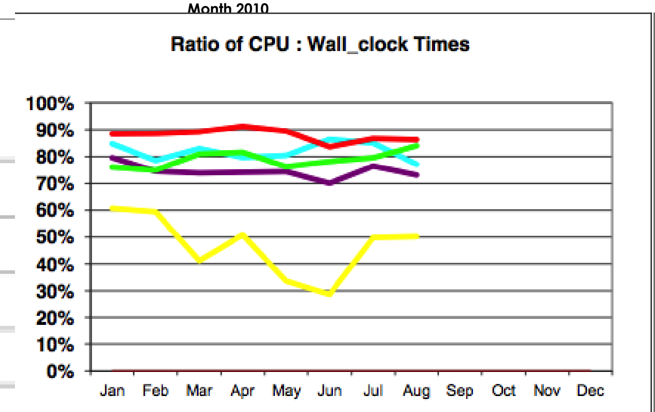

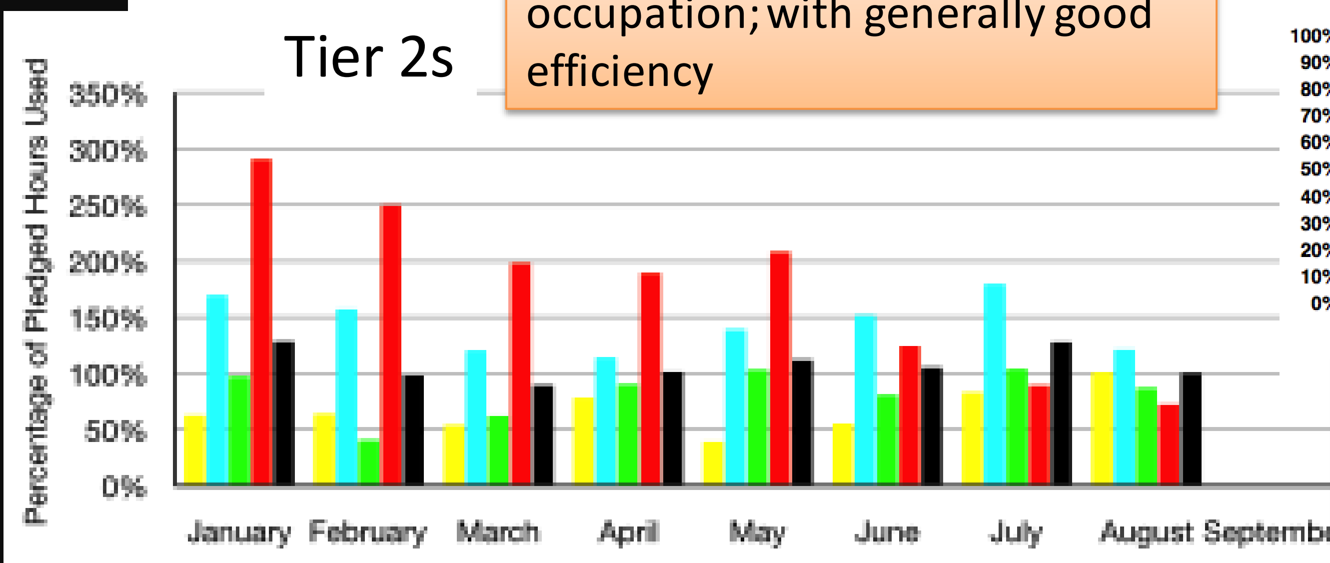

Efficiencies now good; ALICE problem now understood

Experiment computing activities

- Ongoing data taking, processing, analysis
  - Very much as anticipated
  - 2011 looks like a real "nominal year" for computing
- Reprocessing of data foreseen for end of year
- Next slides show a few examples…

LHC performance impact on data volumes

ATLAS

Mean Number of Interactions/Crossing

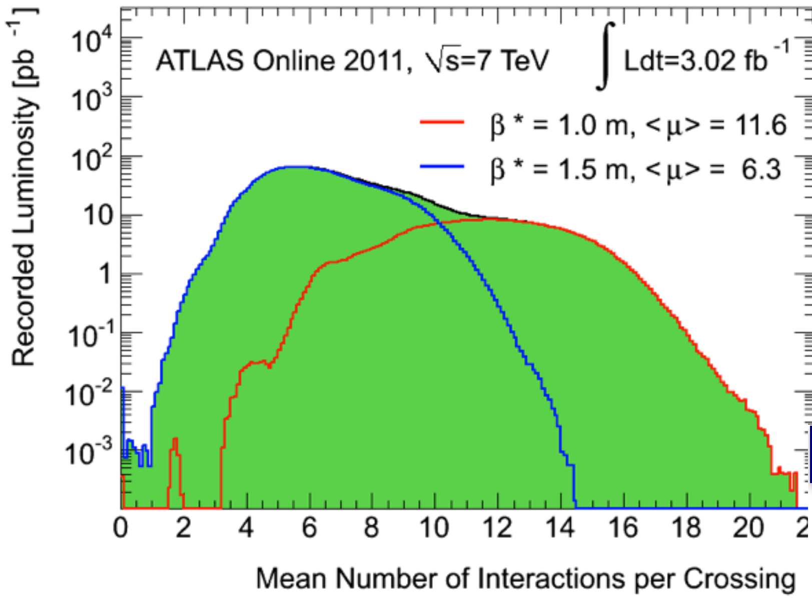

- 965M events recorded since Jan 1
- Two more pp runs in September (27 days) and October (29 days)
  - 30% duty cycle 174M events @120 Hz (~380MB/s), 10% MB and and 90% rare triggers
- PbPb in November (17 days) and December (7 days).
  - 30% duty cycle 200M events @320 Hz (~4GB/s)

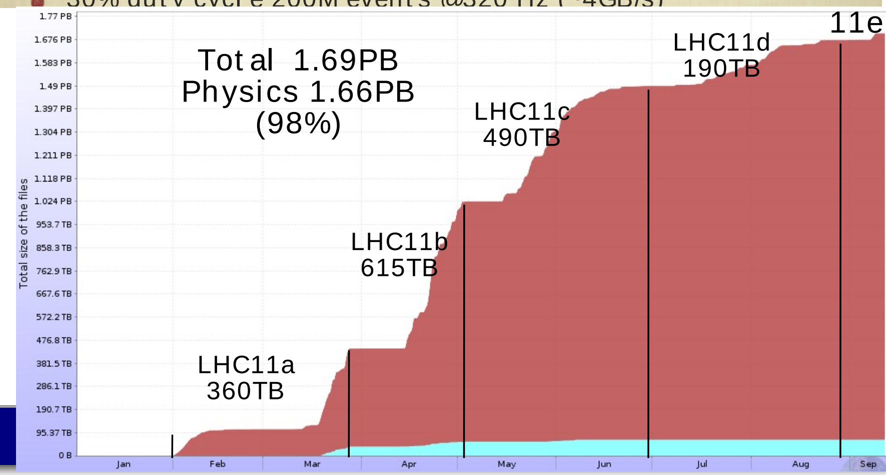

- Machine performance has averaged to about what was planned for
- Months with technical stops are lower, others are higher

CMS

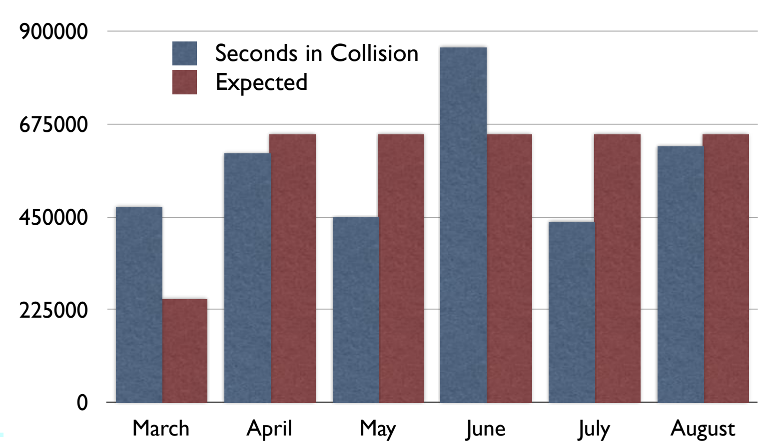

LHCb

Constant luminosity means constant trigger rate, and data volumes proportional to time in colliding beams

- No surprises:
  - 3kHz trigger rate as expected
  - Machine live time about as expected
  - Pileup as expected
- Data volume and processing times as expected

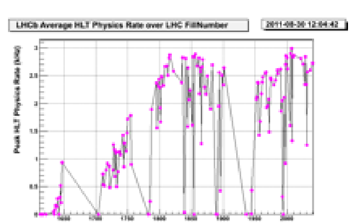

Trigger rates and event sizes

- Event size is very close to estimates
  - MC Reco is larger due to out of time pile-up, which wasn't in the original planning
- In general the RECO time is about 20% higher than anticipated, but this needs to be watched as the pile-up increases

Tier	Observed Size	Expectation
Data RAW	230kB (350kB after TS)	390KB
Data RECO	590kB	530kB
Data AOD	165kB	200KB
MC Reco	970kB	600kB
MC AOD	250kB	265kB

Month	Average Trigger Rate (with overlap)
March	356Hz
April	334Hz
May	393Hz
June	431Hz
July	361Hz
August	380Hz

CMS

- Event sizes are smaller than estimated.
  - Much work went into optimizing sizes after we made our estimates in March 2011
  - This allowed us to run at higher trigger rate than the baseline 200 Hz

ATLAS

Pileup (as shown earlier)
- $\langle \mu \rangle$ = 6 interactions/crossing
- now μ = 15 at start of fills

- Average (Jan.-July) physics trigger rate was 275 Hz
- Typically 400 Hz at beginning of fill (occasionally up to 475)
  - In order to have a uniform data sample for the summer conferences.
- We accommodated this by changes to data distribution, reconstruction algorithms, and event sizes despite pileup
- This is an ongoing physics optimization

Software improvements

Memory profiles

- Simulation/reconstruction – under control (SUCCESS!)
- User analysis – some progress, but still quite high and uneven

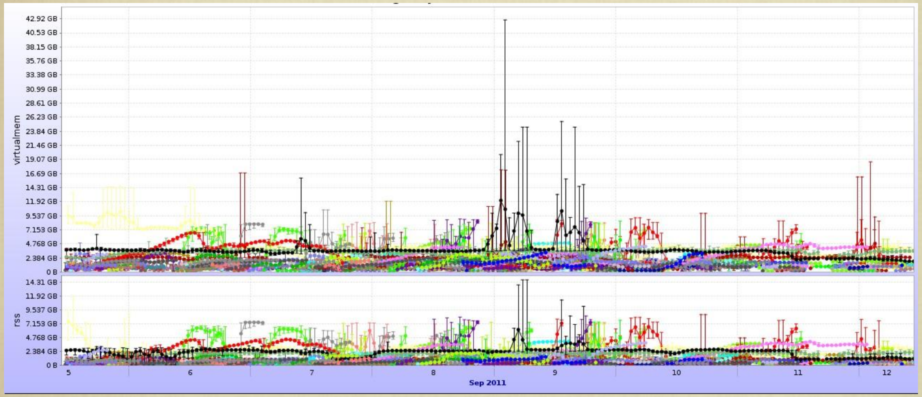

- The spikes still cause troubles at many sites – fast allocations kill the WNs

Grid operation - efficiency

- Efficiency (CPU/Wall) time for ALICE jobs has been 'sliding' over the year

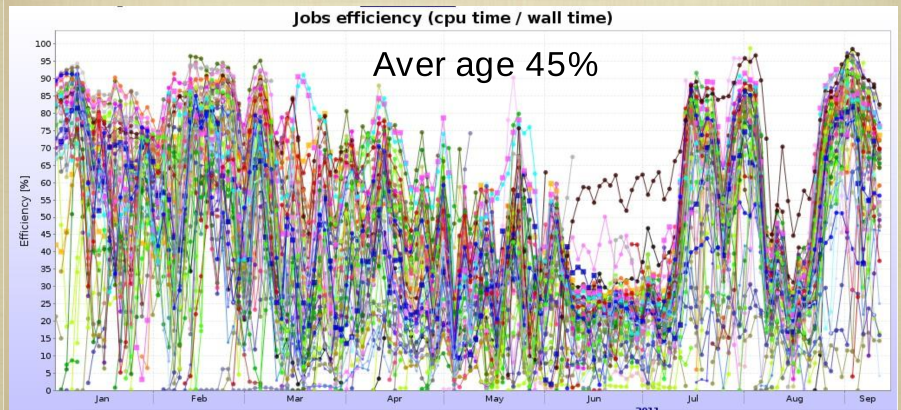

CMSSW_4_2 vs 4_4

- Memory of the application in 4_2 limits us to around 2k processor cores at the Tier-0
  - 4_4 has much better memory performance

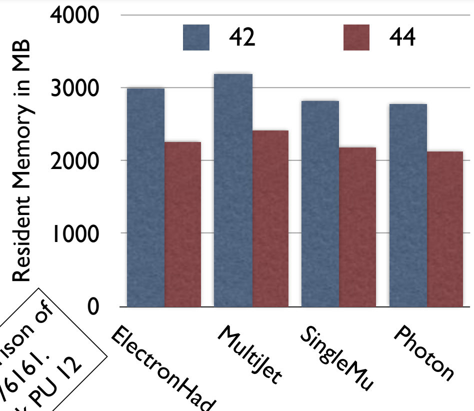

Comparison of 176161. Peak PU 12

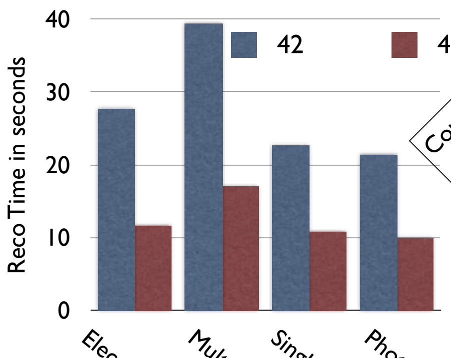

- Reco time estimate in the planning for 16 pile-up events was 11s per event. Possible only in 4_4, we are already at 11s for 12 PU pea

ALICE:

- Doubled number of events / GB with compression in DAQ
- Reduced reconstruction time significantly
- Work hard to fit within available memory (as do several experiments)

CMS:

- Factor 2 improvement in reconstruction time!
- Improved memory footprint – should help with Tier 0 occupancy

Some observations:

- ATLAS:
  - Dynamic data placement: physics datasets copied to Tier 2 only when required; also pre-placement of popular datasets
  - Ntuples often still used for analysis – expected to reduce as datasets get larger: expect consequent increased use of Tier 2 for analysis
- ATLAS and CMS:
  - Move from ESD to AOD for analysis going well (ATLAS have stopped using bulk ESD)
- CMS:
  - After reduced β* Tier 0 became fully occupied – larger events meant more memory & worse utilisation: CERN added RAM in machines + improved sw
- LHCb:
  - Fully utilize Tier 1s: OK for data taking, but could not run re-processing in parallel: started to commission Tier 2s for reprocessing

Resource request adjustments for 2012

- ATLAS:
  - Double Tier 0 CPU request (73→111 for CERN total)
    - In order to maintain current trigger rate
- CMS:
  - Increase CERN disk by 2 PB in order to effectively use available CPU for analysis
- ALICE:
  - Large CERN disk request (as in Spring) to manage increased event rate and need to keep data on disk
- LHCb:
  - Missing 2 PB disk (globally) – several Tier 1s have offered space

Site reliabilities …

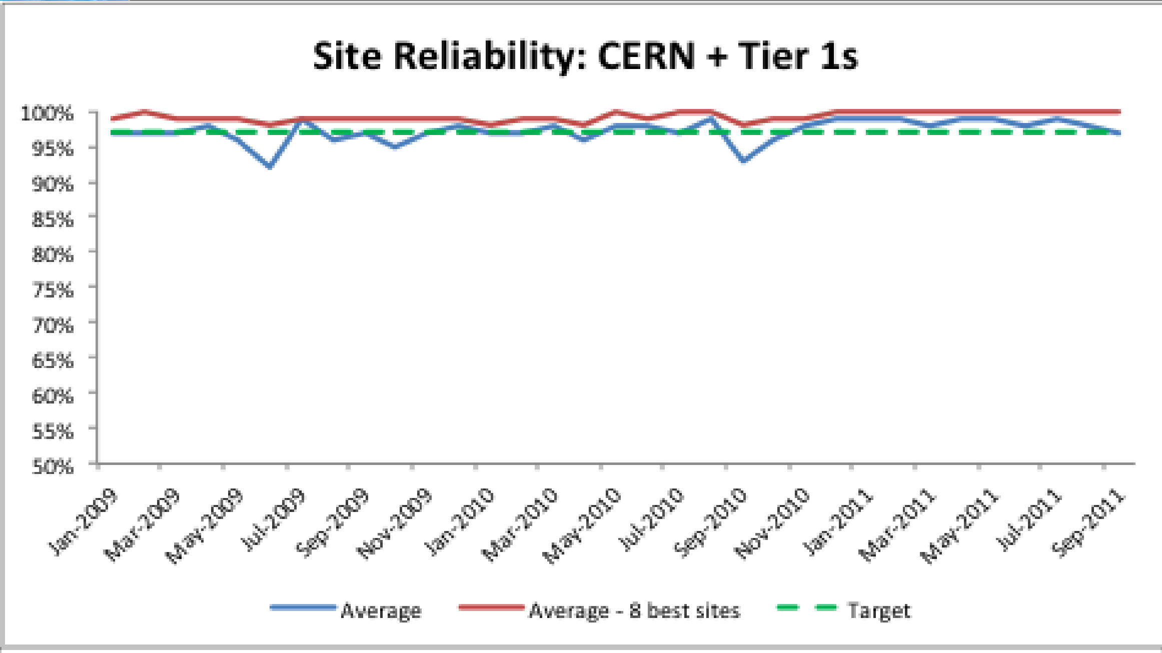

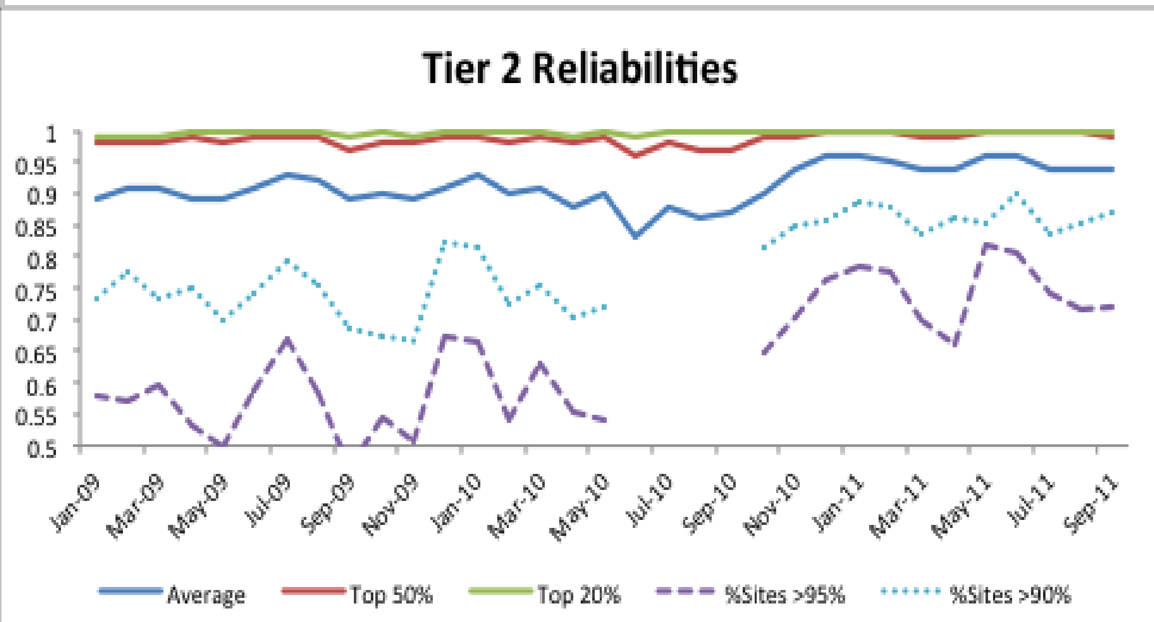

Mostly this is not an issue now;

More than 75% of sites have reliabilities better than 95%

Regular measurement of availability by ops and experiments are rather positive

… Service incidents

- Decrease in SIR’s wrt 2010 while greater service load
- Correlation between service changes and problems
- Response time well within targets
- More problems resolved in 24 hrs wrt 2010
  - Fraction of incidents take longer than 96 hrs – complex problems

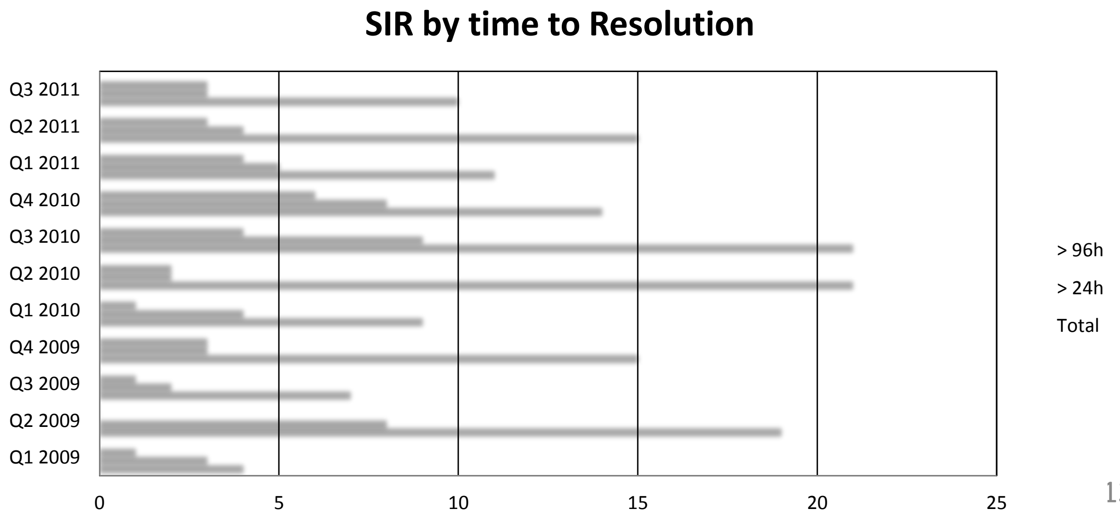

Service incidents

- Main sources now:
  - Infrastructure/fabric services – typically power/cooling related or hardware failure: constant level
  - Storage services – at all sites and critical at Tier 0/1
  - Database services – mainly at Tier 0/1

Incident rate has decreased, despite heavy load, although complex problems can still take long time to resolve

✔ Operations load is now viewed as reasonable and sustainable by sites

Status of Tier 0

- Many responses to request for proposals
- (Almost) all have been followed up with site visits and extended discussions
- Tender documents prepared (specs and SLA) and sent out
- …
- Anticipated timescales:
  - Tests in 2013 (optimistically already in 2012), production in 2014

LHCONE Status

- **Intention was to execute a prototype proof of concept**
  - **Based on the concept of Open Lightpath Exchanges.**
  - **Switched core, routed edge.**
  - **Recognise individual contraints**
    - Allow everyone to use existing equipment.
    - Use existing standards.
- **This has resulted in an interesting and working proto-infrastructure**

Overview (from Dante)

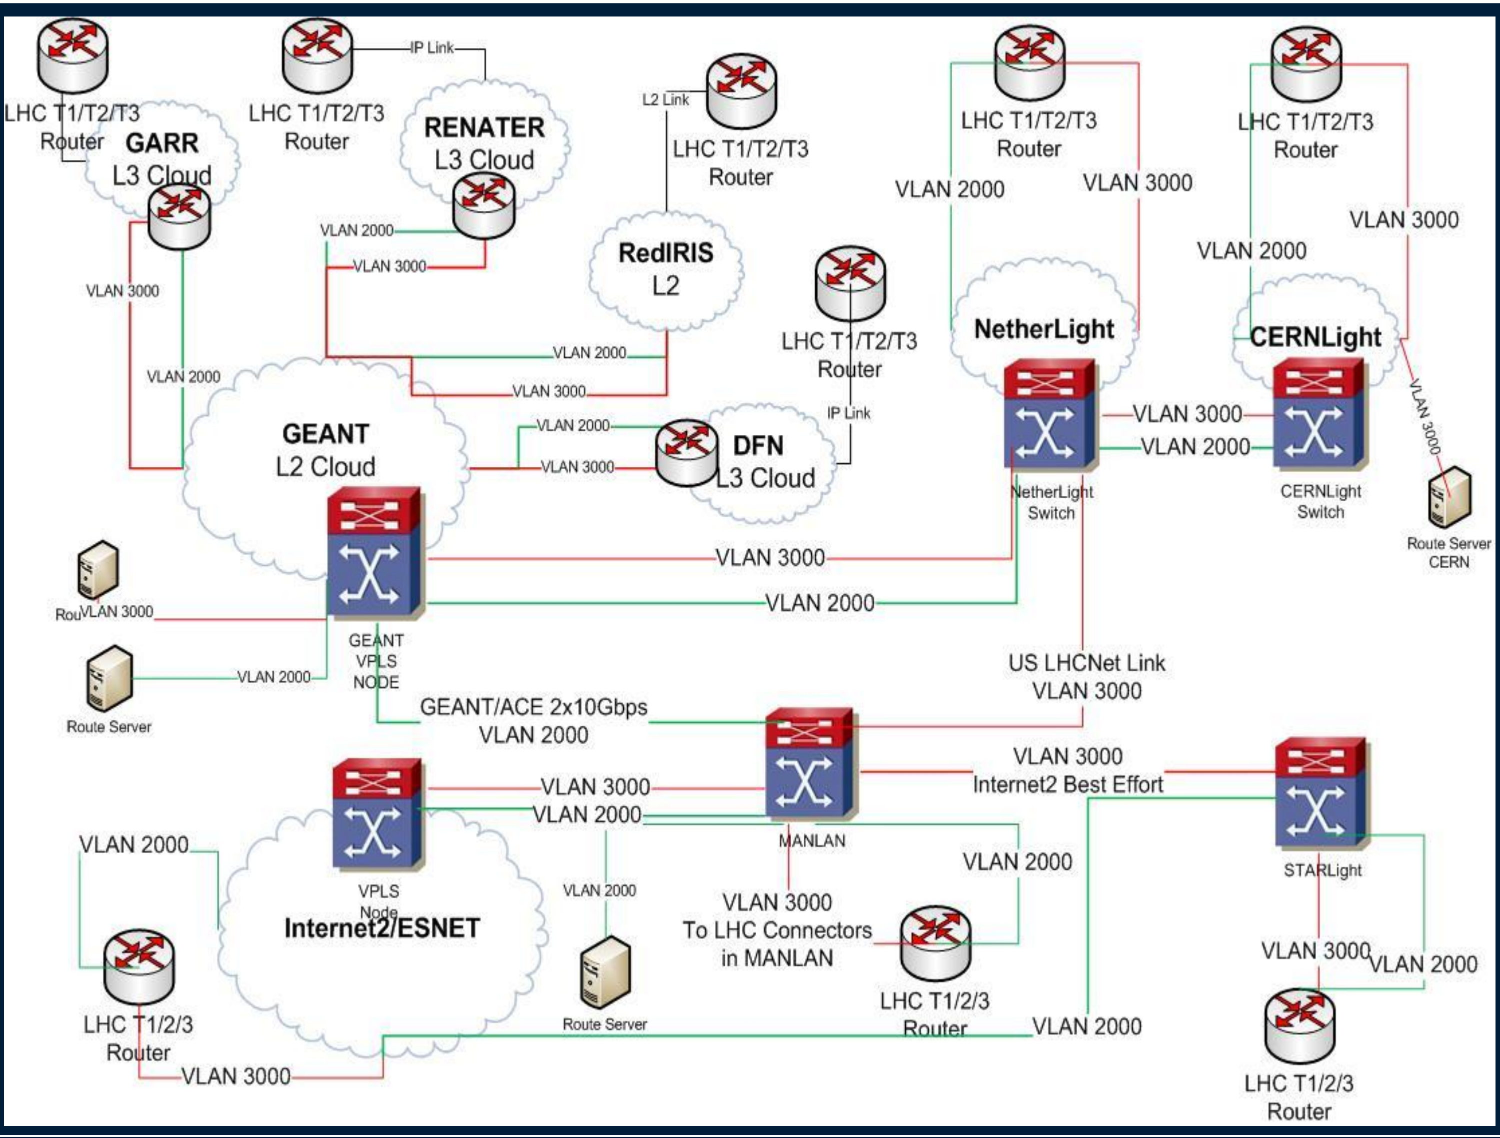

LHCONE Status

- **This has resulted in an interesting and working proto-infrastructure**
  - **But scalability and manageability are doubtful with current technology.**
- **This has resulted in a fork in the road**
  - **Transform this into a safe, but constrained solution today**
    - Are the requirements from the R&E networks still valid?
  - **Slow down the move to a production infrastructure to investigate emerging solutions**
    - New standards for L2 networking.

Technical evolution: Background

Consider that:

- Computing models have evolved
- Far better understanding of requirements now than 10 years ago
  - Even evolved since large scale challenges
- Experiments have developed various workarounds to manage weaknesses in middleware
- Pilot jobs and central task queues (almost) ubiquitous
- Operational effort often too high; lots of services were not designed for redundancy, fail-over, etc.
- Technology evolves rapidly, rest of world also does (large scale) distributed computing – don't need entirely home grown solutions
- Must be concerned about long term support and where it will come from

Strategy

WLCG must have an agreed, clear, and documented vision for the future; to:

- Better communicate needs to EMI/EGI, OSG,…
- Be able to improve our middleware stack to address the concerns
- Attempt to re-build common solutions where possible
  - Between experiments and between grids
- Take into account lessons learned (functional, operational, deployment, management…)
- Understand the long term support needs
- Focus our efforts where we *must* (e.g. data management), use off-the-shelf solutions where possible
- Must balance the needs of the experiments and the sites

WLCG
Worldwide LHC Computing Grid
LCG

Strategy: Working groups

- Agreed at collaboration meeting: working groups set up to address key topics and produce strategy for the future
- Groups:
  - Data Management
  - Storage Management
  - Workload Management
  - Databases
  - Security
  - Operations and tools
- Timescale: 1^{st} draft Jan/Feb 2012

Conclusions

- Grid operations have continued smoothly over 2011, no major issues
- Experiments make good progress in data processing and analysis
- Tier 1 and Tier 2 resources fully utilized
- Tier 0 fully utilized during peak times, could be better used overall
- Additional resource requests have been made for 2012
  - Main implications are at Tier 0
- Planning has started to produce strategy document for evolution of the WLCG infrastructure

WLCG Worldwide LHC Computing Grid

LCG