LOI BENCHMARKING PRODUCTION ON THE LCG

Marcel Stanitzki, Jan Strube
STFC Rutherford Appleton Laboratory

The LHC Computing Grid

- Will be able to process millions of files for you
- Just look like a normal batch queue
- Will host the thousands of Terabytes of data produced by the LHC experiments
- Will serve several thousand scientists

The LoI Samples

- SM Backgrounds
  - 250 and 500 GeV
- Various Signal samples
  - ZH & ZHH inclusive
  - Charginos and Neutralinos, sbottoms
  - Top pairs
  - $H \rightarrow cc, \mu\mu$
- Beam pairs

Getting Resources

- Resources are allocated by regional experiment boards
- For the RAL TIER1: GridPP
- Need to request
  - CPU time
  - GRID Storage (CASTOR)
- We needed to cut some red tape here
  - LHC has priority usually
  - But LHC is broken (golden window of opportunity)

The CASTOR ..

- A mass-storage system developed by CERN
- Basically a huge tape library
- Accessible as a filesystem and with GRID tools
- The theory so far
- In practice
  - Only works reliably at CERN
  - On other sites problematic
  - Needs a lot of hand-holding

Production Stages

- Simulation (STDHEP → LCIO)
  - Using SLIC
- Reconstruction (LCIO→ LCIO)
  - org.lcsim
- LeptonID (LCIO→ LCIO)
  - org.lcsim
- Vertexing (LCIO→ LCIO)
  - MarlinReco

See https://confluence.slac.stanford.edu/display/~jstrube/LOI+Analyses+Bookkeeping

The big board ...

Sample name	stdhep	SLiC	reco	leptonID	MarlinReco
slac_susy_point5_delMneu2_0p5	1277	1277	1277	1277	-
ffhh_ghhh1p00	194	194	194	194	194
SM_Sample 250 GeV	936	936	7065	7879	7879
Bhabha	6317	6317	6296	6296	-
slac_susy_point5_delMneu1_0p5	1205	1205	1205	1205	-
H mu mu	102	102	102	102	-
sbottom MSB230_MNE220	198	198	198	198	198
ZH generic	-	-	215	215	215
sbottom MSB240_MNE210	194	194	194	194	194
ffhh_ghhh1p25	220	220	220	220	220
sbottom MSB230_MNE210	199	199	199	199	199
sbottom MSB240_MNE220	199	199	199	199	199
eeh mumuh	602	602	602	602	602
beam backgrounds	26454	26454	-	-	-
slac_desy_point5_map_v2	107	107	107	107	-
sixfermion mtop 173.5 GeV	-	-	1072	1072	1072
slac_susy_point5_delMch1_0p5	1294	1294	1294	1294	-
sixfermion mtop 174.0 GeV	-	-	1108	1108	1108
SM_Sample 500 GeV (new)	575	575	6622	7182	7182
slac_susy_point5_delM_0p0	1264	1264	1264	1264	-
Files (1000 events each)	**41337**	**41337**	**29433**	**30807**	**19262**

Reality strikes ...

- Will be able to process millions of files
  - As long as you don’t submit more than 50 jobs at a time
- Just look like a normal batch queue
  - But it takes 4-6 weeks before you can actually start
- Will host the thousands of Terabytes of data produced by the LHC experiments
  - But if your logs grow larger than a few hundred MB, you grind some sites to a halt
- Will serve several thousand scientists
  - As long as don't you work on more than one experiment

What worked well ...

- CPU time / wall clock time / memory limitations
- disk space (!)
- org.lcsim and Java on the grid
- sending jobs to MarlinReco
- RAL, DESY, IFH sites

What did not work so well

- Disk access by multiple users / processes
- Copying files (in and out of CASTOR)
- File validation of various stages / error checking
- Job monitoring / logging / GANGA
- Most other grid sites are troublesome
  - Reliability
  - Broken setups

Other comments

- The GRID is designed towards LHC
  - A lot of dedicated manpower from the LHC experiments as well
- Non LHC experiments …
  - You are on your own (almost)
- Job failure rate is appalling
  - Numbers from LHC: Need to submit each job 1.8 times
  - We had similar experiences …
  - Mostly grid-specific failures !

Some plots SM 250

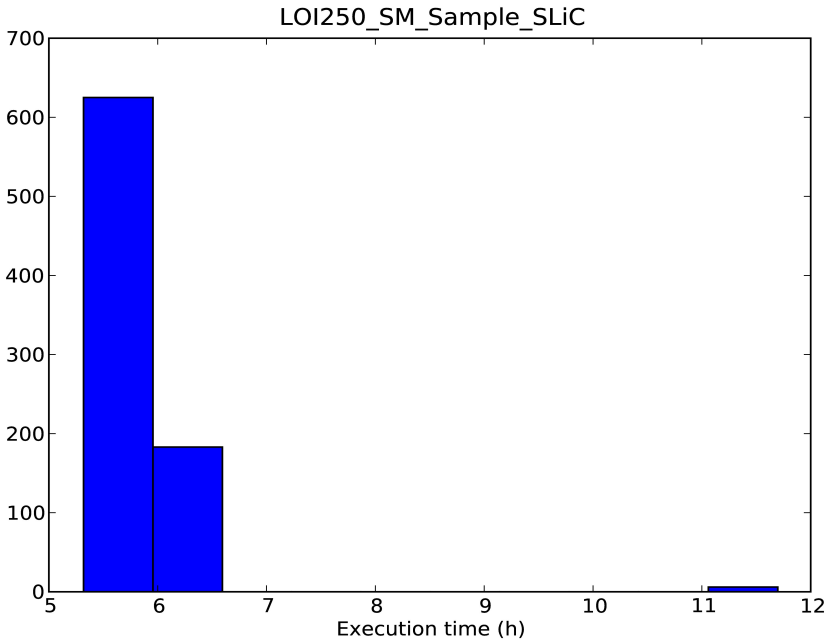

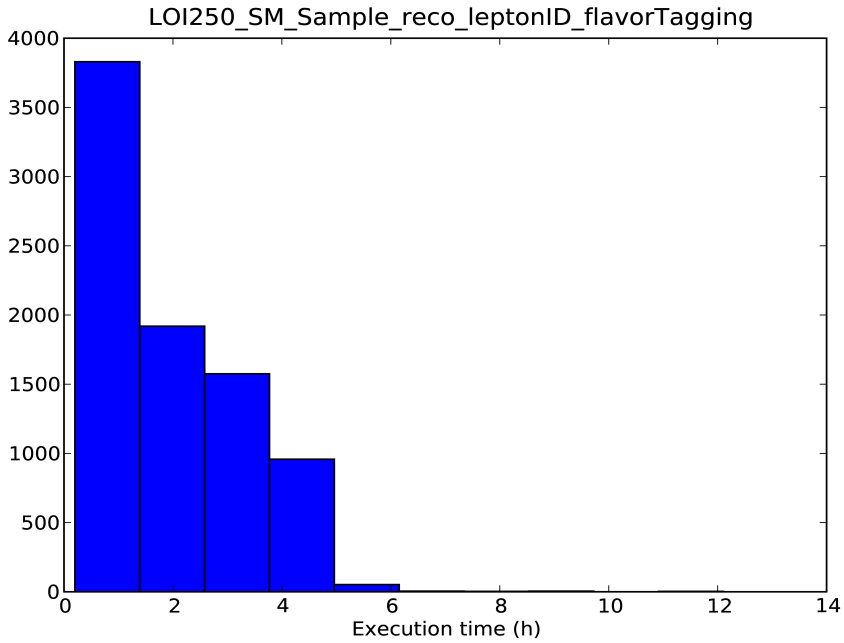

Beam background

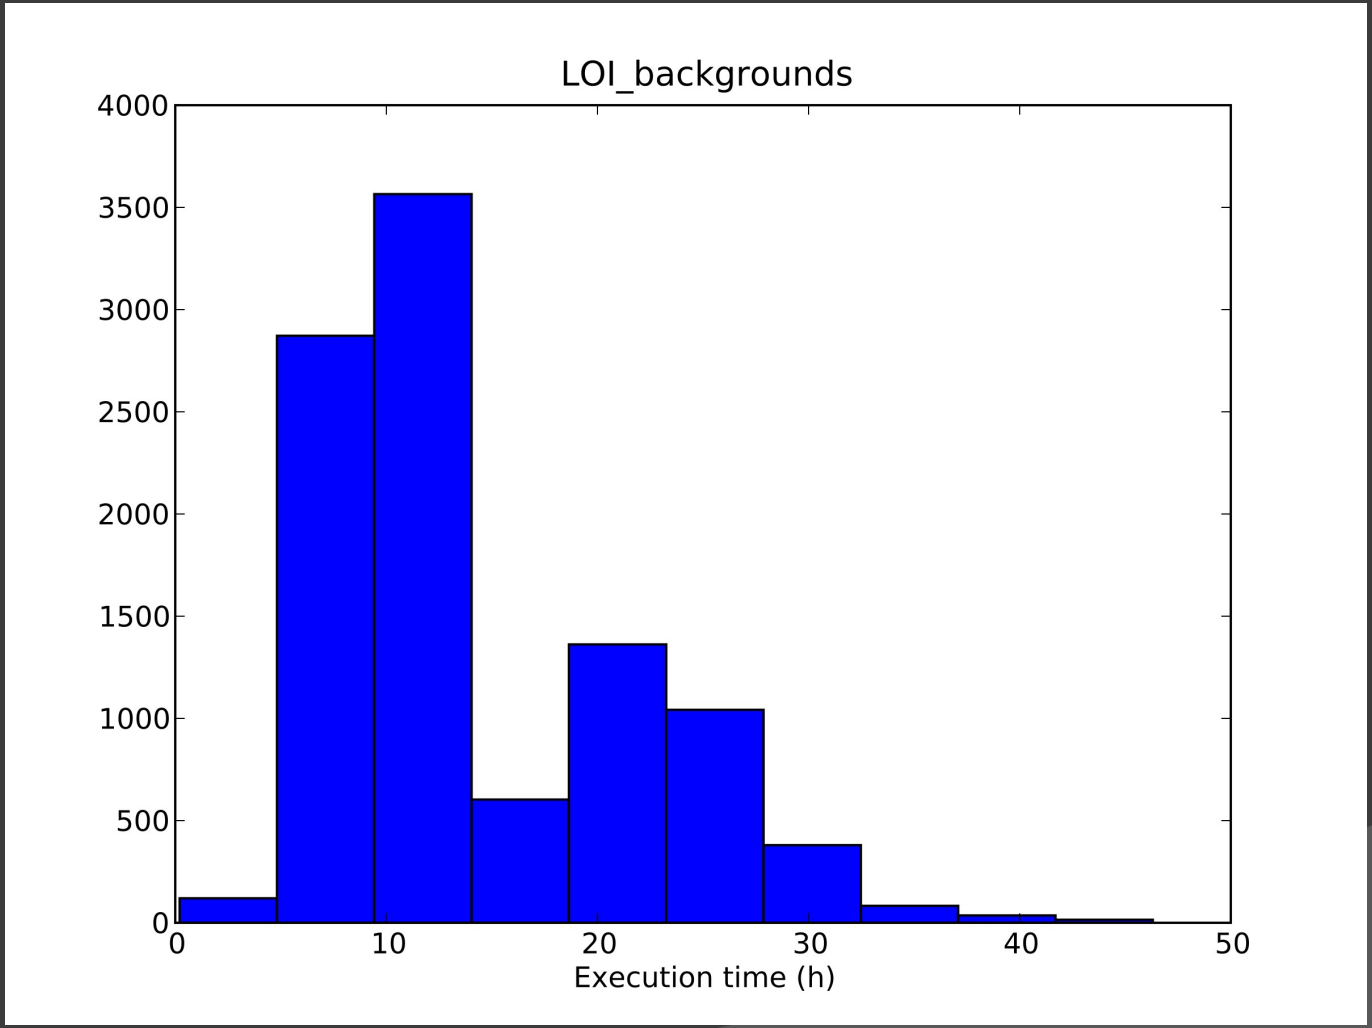

Putting it together

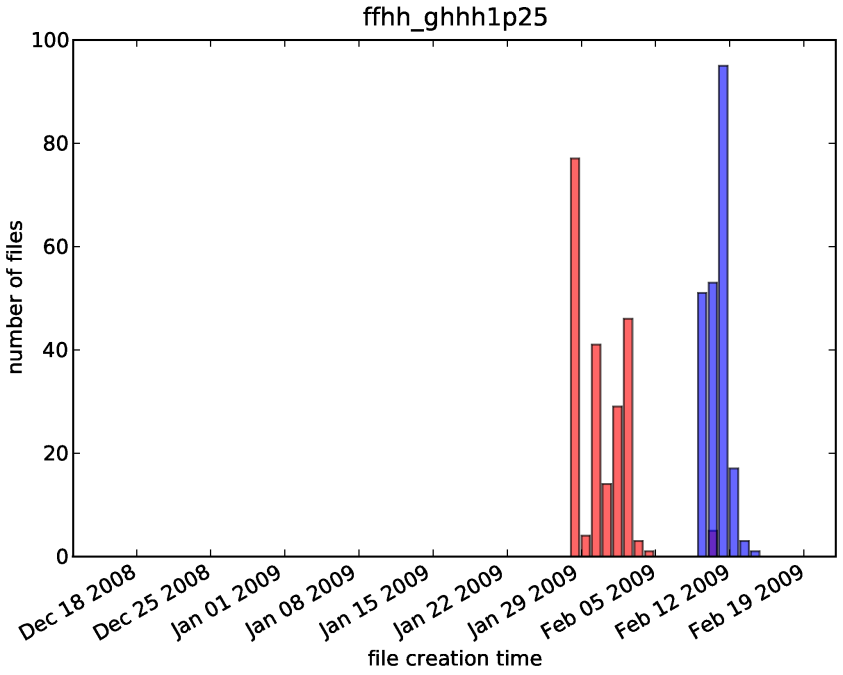

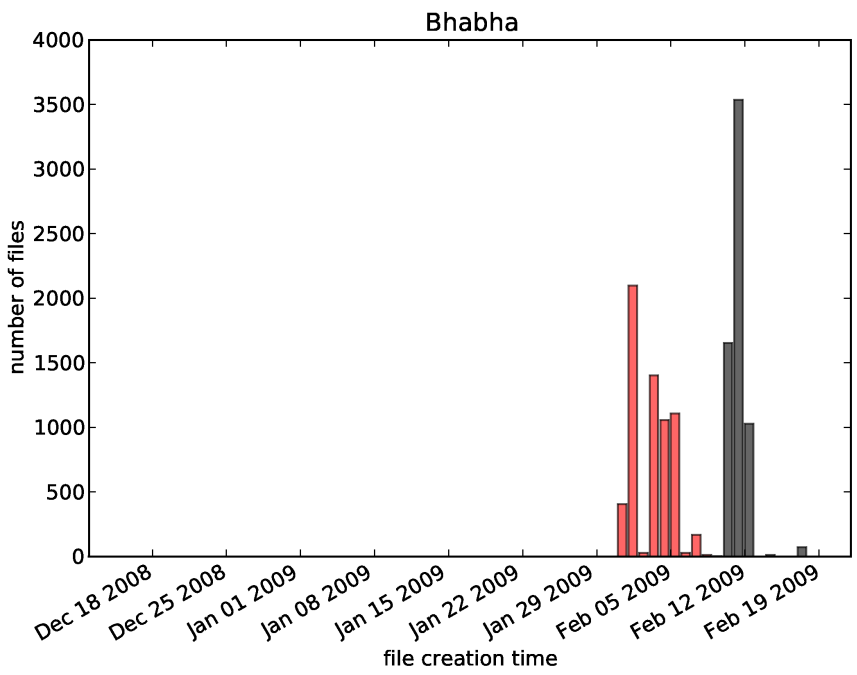

Reco + leptonID

SLiC

Reco + leptonID + MarlinReco

GRID CPU time

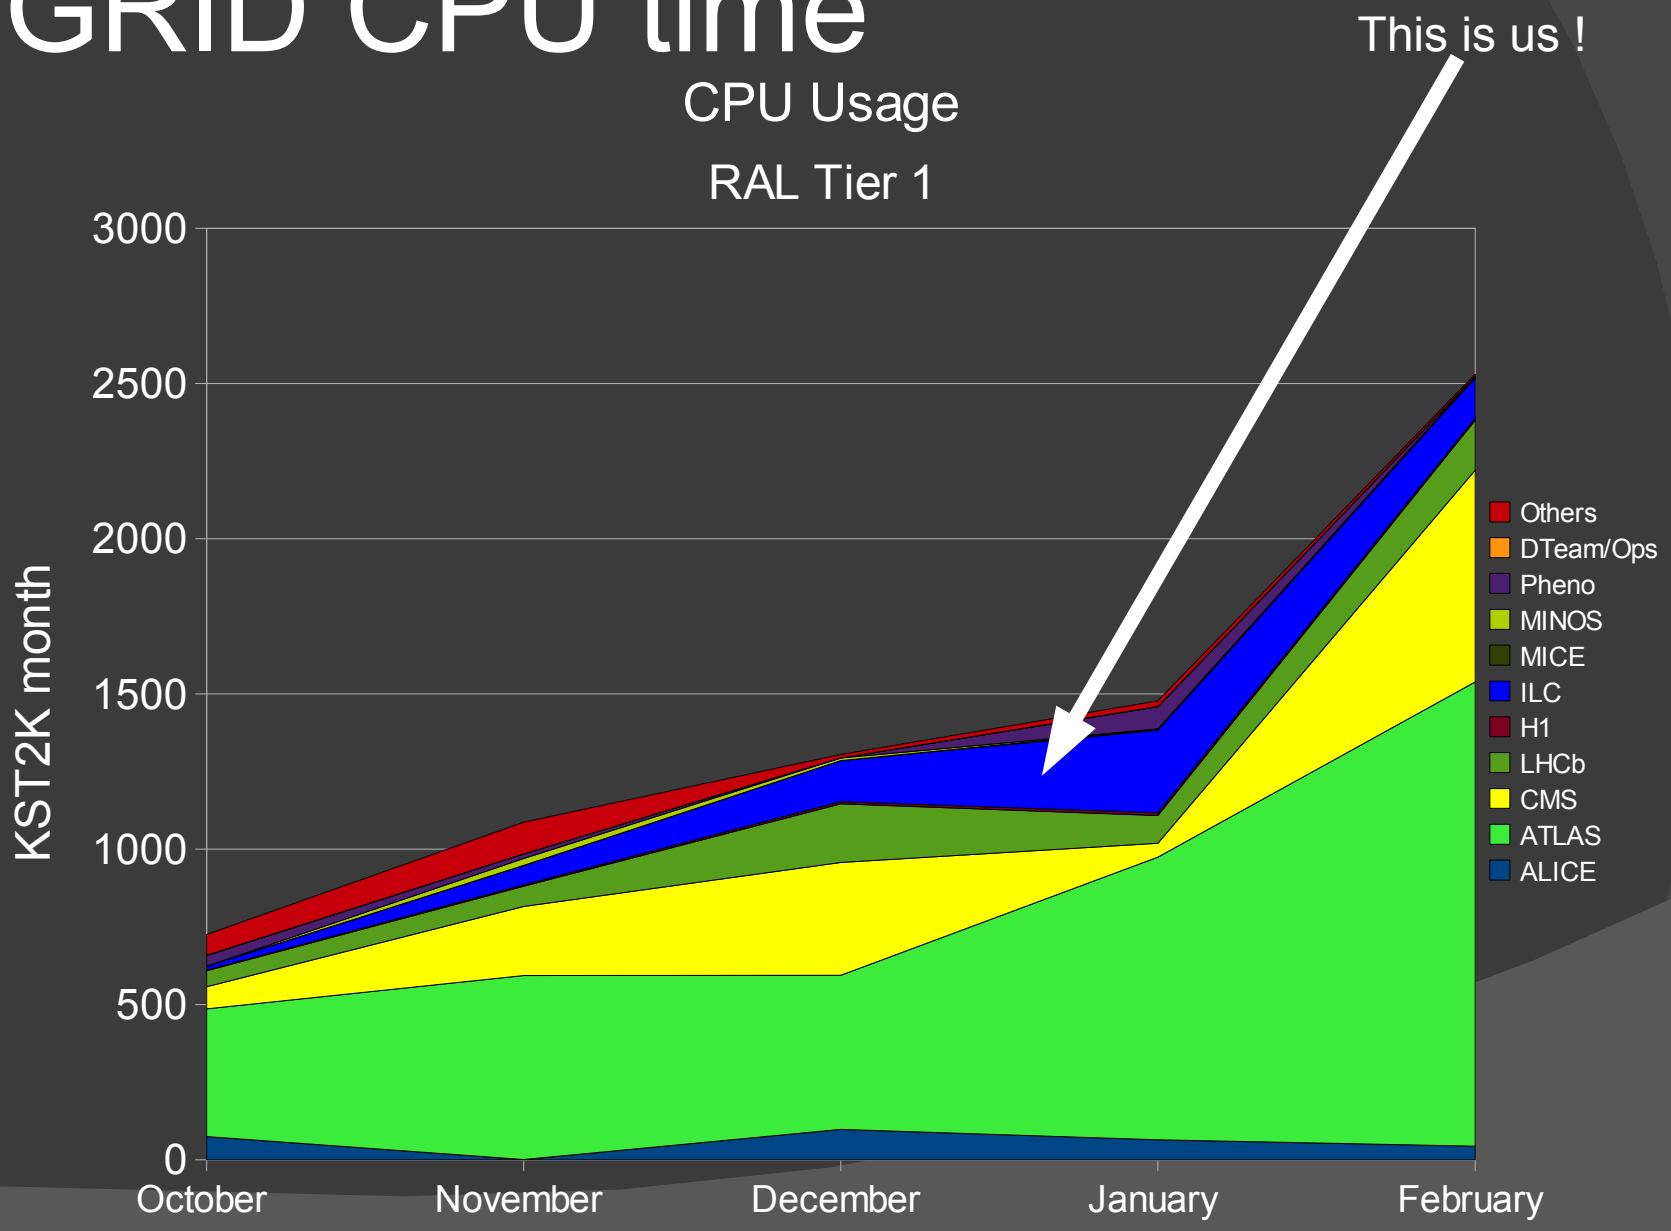

Summary

- We're almost done (just a few things missing)
- A big thanks to Jan for being the production coordinator
- Thanks to Glenn Patrick (GridPP Allocation Chair)
  - For getting us CPU and storage on the RAL Tier1
- A big thanks to the GridPP team for supporting us
  - Gareth Smith, Andrew Sansum, Shaun De Witt and the other members of the RAL Tier1 Team