CMS Report at WLCG LHCC Review March 2016

Daniele Bonacorsi, David Lange

March 1, 2016

Outline

- Resource usage
- 2016 plans
- HL-LHC activities

Some Recent achievements in CMS software/computing area

- CMS completed our **end-of-year reprocessing**
  - We reprocessed both data and Monte Carlo using updated software to address changes that accumulated during the 2015 run
  - This production was the first use of our multi-threaded application on the GRID.

- CMS has achieved 200k simultaneously running jobs using **global pool** and **commercial** resources

- CMS **Amazon project** achieved its goal of more than 50k simultaneous running jobs
  - Demonstrated at the scale of the sum of CMS Tier-2s.
  - Monte Carlo output in use for Moriond analysis

- CMS is preparing for **2016 operations**
  - CMS digitization and pileup simulation software now multithread efficient
  - Now working to move the largest Tier-2 centers to use multi-core pilots

Stable operations across CMS computing centers

2015 T1 CPU

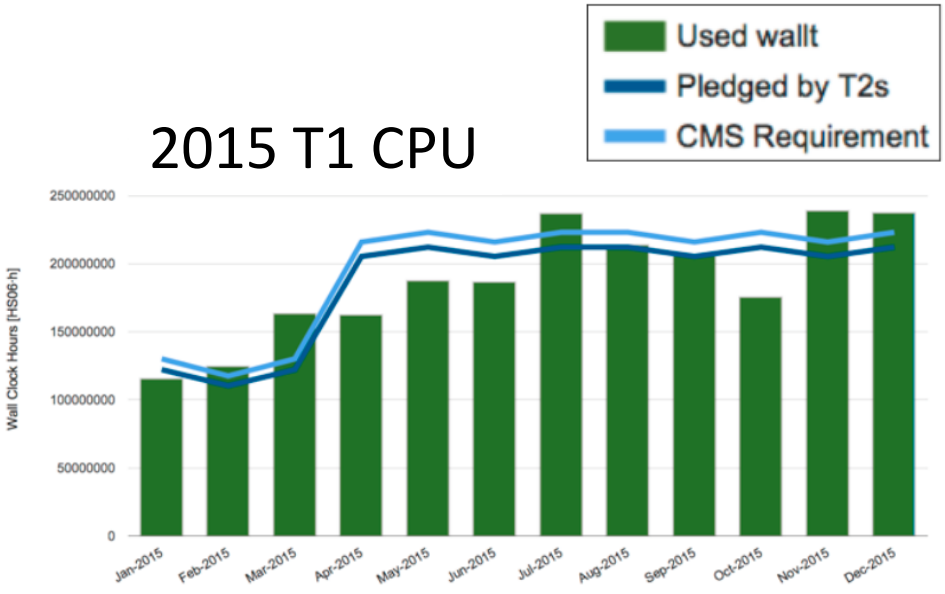

2015 T2 CPU

We have just completed the major end-of-year re-reconstruction for 2015 data and MC

Building on LS1 improvements to exploit new resources

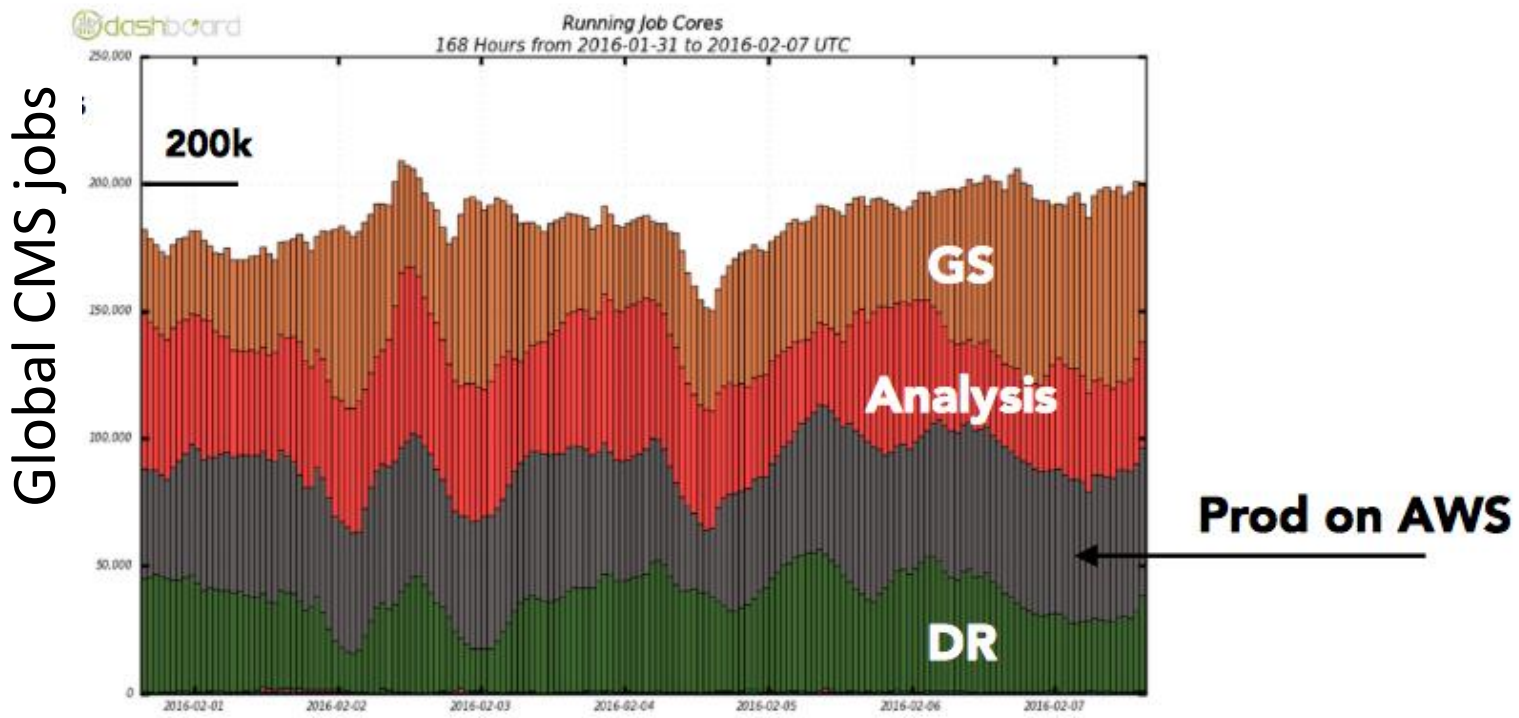

CMS reached 200k running jobs including commercial cloud and HLT resources

HLT farm now in routine use for production (both SIM and our high-I/O DIGI-RECO)

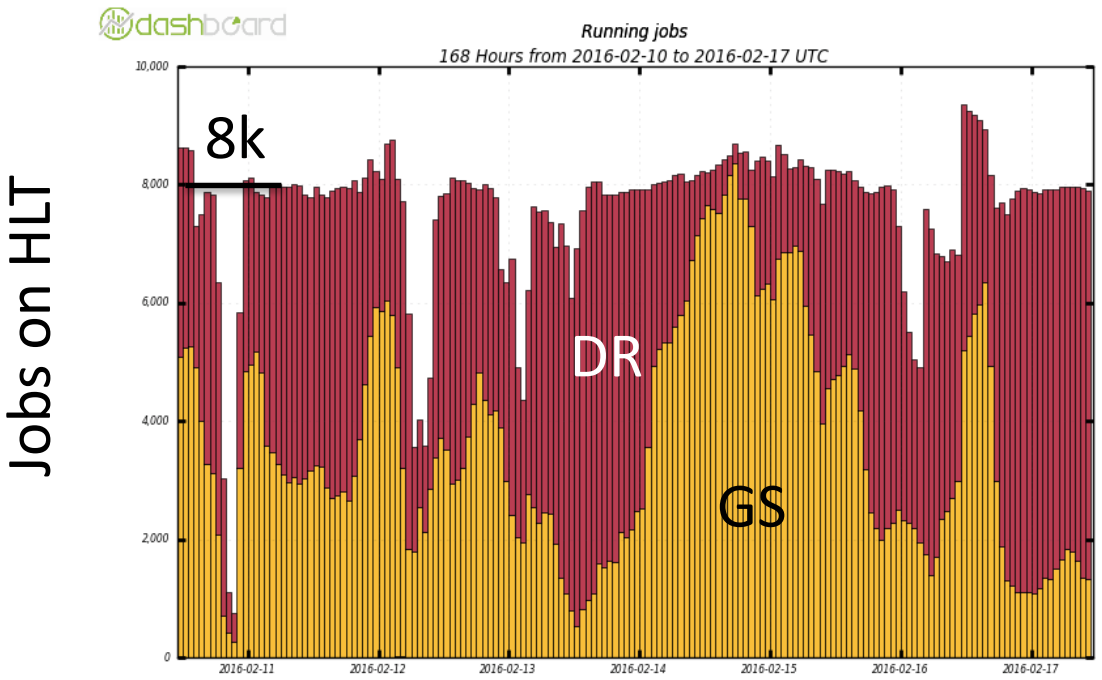

Given the planned major CPU upgrade in April, the HLT will be an especially important resource for CMS

Analysis usage at Run 1 levels and growing to complete Moriond results

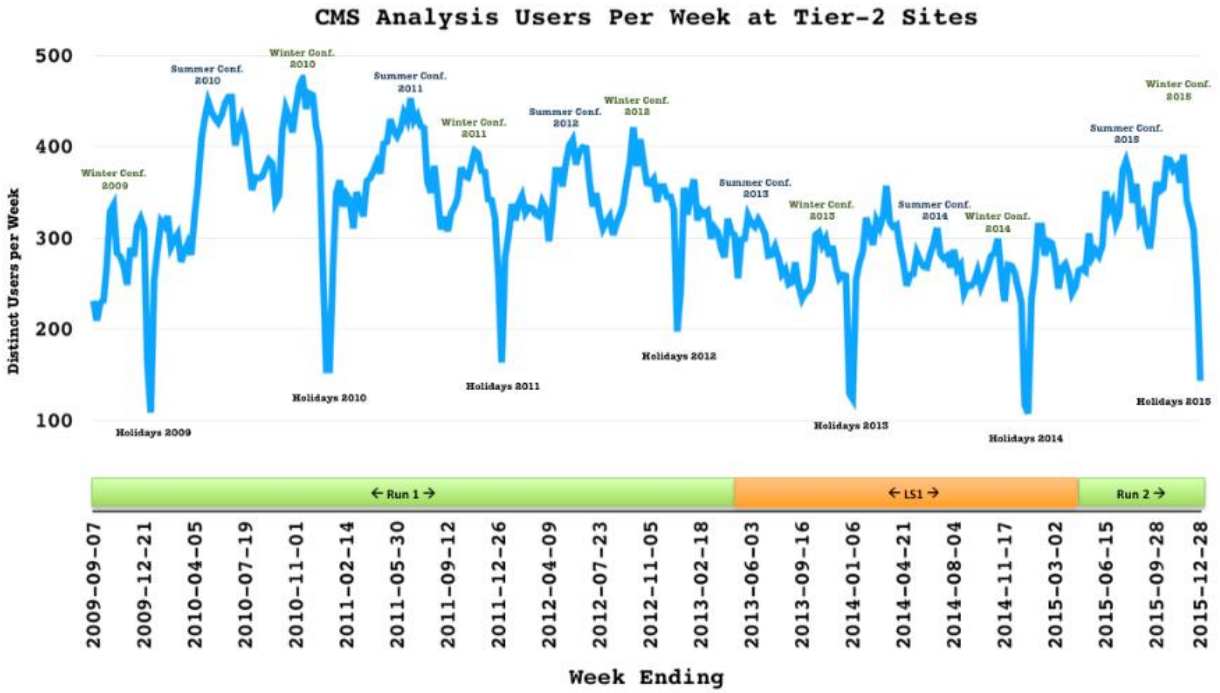

- A big increase in analysis work during January and February is the biggest test so far of our LS1 operational changes (and specifically to CRAB).
- MiniAOD adoption continues. We made 4 versions of it this year.

End of year data re-reconstruction was first use of CMS multi-threaded application beyond Tier-0

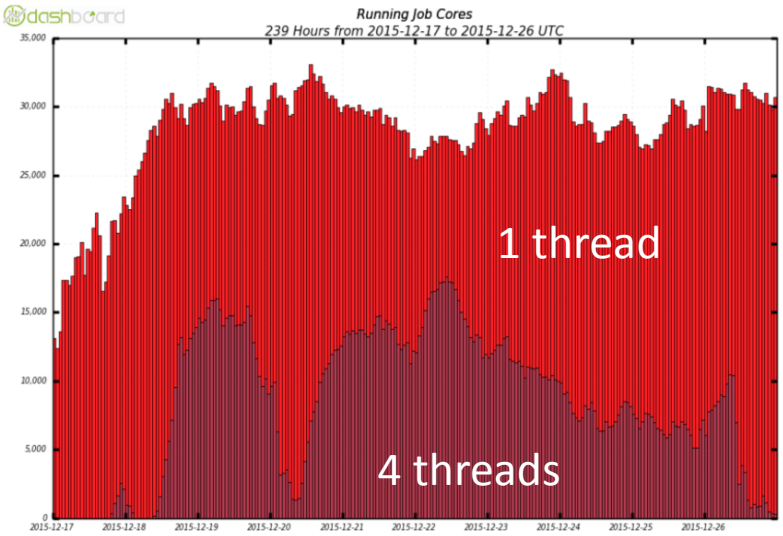

~50% of CMS Tier-1 resources used for re-reco. over holiday period

- For 2016 we expect DIGI-RECO MC and data processing to be multi-threaded except on sites with single core pilots
  - GEN-SIM application is also ready, but minimal production planned until 2017 sample production begins late in the year
- This change will help in a number of ways. These include:
  - Less stress on workflow management system (fewer simultaneously running jobs)
  - Better use of CMS CPUs (all workflows under 2GB/core)

Preparation for 2016

- "Global runs" have already started. We are using latest CMSSW software on both Tier-0 and HLT

- 2016 Monte Carlo production starting in ~1 month
  - New software, in particular for our upgraded L1 trigger
  - We will re-use the detector simulation from 2015 (GEN-SIM samples). The DIGI-RECO processing is planned to use 4 threads on sites that have multicore pilots

- Given end-of-year processing work, disks are nearly full of "last copy" data (not automatically cleaned)
  - This reduces the effectiveness of our dynamic data management system
  - We know some of our data is obsolete: We have organized a cleanup campaign (aiming for 10 PB reduction of samples in disk cache)

CMS software+computing towards HL-LHC

- CMS expects to need very high trigger output rates to deal with increased event complexity of HL-LHC operations

- At high pileup, DIGI-RECO processing is by far the biggest CPU need
  - DIGI-RECO is already a large fraction of CMS processing. Its processing time per event scales worse than linearly for DIGI-RECO workflow.
  - The SIM (Geant) processing time per event is essentially unchanged with pileup as we simulate pileup events separately from the hard scatter

Size of the processing problem if we continue as we continue with our current model

- Factoring in the trigger rate and taking a weighed average of the data and simulation tasks we see the computing challenge is 65-200 times worse than Run2 (2016-like conditions)

Detector	HLT output rate (kHz)	Total
Phase 1	1	3
Phase-II (140)	5	65
Phase-II (200)	7.5	200

Scale of computing resource needs relative to Run 2 including the increase in projected HLT output rate

(More detailed estimates in our 2015 Technical Proposal)
https://cds.cern.ch/record/2020886?ln=en

Even given technology evolution, we have big deficits to make up via R&D

We considered the simplest scenario.

Assumptions:

1. Flat budgets for computing
2. Resource capability per CHF will continue to scale as they have in the past

Example results:

- Anticipating a factor of 8 in CPU/CHF improvements, we would have a deficit of a factor of **6-30** to regain from application improvements (etc.)
- Anticipating a factor 6 in storage/CHF improvements, we would have a deficit of **4-5** in storage

Example R&D solutions under investigation

1. **Optimizing on new architectures and technology**
   - Porting CMSSW to new architectures as they are available via TechLab or other resources
   - Pushing development of multithreaded framework and algorithms for these architectures
2. **Code algorithmic improvements**
   - We achieved large reductions in processing time during LS1 (2x for simulation, 3x for reconstruction). A small team is constantly looking at how reconstruction time per event can be reduced.
   - Exploring alternative approaches: Premixing is an example of how we can make a big reduction in both I/O (to improve CPU efficiency) and overall CPU needed for DIGI-RECO
3. **Agility to take advantage of resource diversity**
   - Commercial cloud projects
   - Use of specialized computer centers
4. **Data reduction and selection techniques**
   - A number on-going projects evaluating analysis, production and computing tools using approaches such as Spark and Hadoop

Conclusion

- 2015:
  - We just completed a major reprocessing pass. We were very successful in delivering large samples quickly for Moriond analysts
- 2016:
  - We are in the final stages of preparation for the 2016 Monte Carlo production campaigns.
  - We expect CMS resources to be busy given the LHC+CMS plans.
  - In addition to normal operations, CMS has a major detector change planned for the end-of-year technical stop. Software development for the pixel and hadron calorimeter changes are ramping up
- HL-LHC:
  - Large challenges that need R&D to address between now and HL-LHC. Challenges are not insurmountable given track record of resource optimization in CMS sites and workflows
  - The landscape of resources on the 10 year time horizon is a big uncertainty. R&D directions try to cover the range of possibilities