The 2007 Europhysics Conference on High Energy Physics IOP Publishing
Journal of Physics: Conference Series **110** (2008) 092003 doi:10.1088/1742-6596/110/9/092003

CMS Offline and Computing Preparation for Data Taking

Tommaso Boccali

INFN Sezione di Pisa, Largo Pontecorvo 3, 56127 Pisa, ITALY

E-mail: `tommaso.boccali@cern.ch`

Abstract. The LHC accelerator in Geneva is expected to start colliding proton-proton beams at an energy of 14 TeV by Spring 2008. The CMS Collaboration is finalizing the construction not only of the apparatus, but also of the offline and computing infrastructure, in order to efficiently analyze the first collected data. Offline operations during the pilot run will include reconstruction at the T0, re-reconstruction with updated calibrations and alignment, sample skimming and analysis tasks. On the Computing side, this implies an efficient infrastructure able to deal with 300 Hz of DAQ output, and to move in an organized way the multi TB samples from T0 to the T2s, where analyzes will take place using GRID-like facilities. The Offline and Computing infrastructure has already been tested to a 25% scale during the challenge CSA06; a 50% challenge, CSA07, is going to happen during the summer.

1. Introduction

The Commissioning of the CMS detector includes not only the preparation and the verification of the CMS detector itself, but has also to insure a proper functionality of the physics analyses at the start up. Given the complexity of the computing system and the volumes of data involved, and the difficulties of event reconstruction in such a harsh environment, a dedicated commissioning is being performed for Offline and Computing infrastructures. In the following we will try to clarify which are Computing/Offline goals for CMS, the resources we are putting into place, and the strategy to be ready for the first collisions.

2. Computing and Offline system scale

The LHC collider is expected to provide, after a short commissioning, proton-proton interactions at a rate of 40 MHz. A first Trigger Level1 (*L1*) “accept” brings this down to 100 kHz using ad-hoc hardware; the rest of the trigger system (*HLT*) is instead running on commercial PCs, running the same C++ code as the one used for Offline event reconstruction. This is expected to bring down the rate to disk to 100-300 Hz, which translates to 200-600 MB/s to be saved. So much for the data taking, but the Offline and Analisys part is even more challenging: the data from the detector must reach computing centers for reconstruction, to be compared with analogous simulated data, and finally used for the analysis. To handle such a complicated work flow, CMS embraced long ago the MONARC[1] model, in which data centers of different size and complexity (the *Tiers*) contribute to the global effort. We distinguish between the *Tier0*, a few *Tier1*s and some tens of *Tier3*s, with functionalities explained in table 1.

Table 1. CMS *TierX* goals and integrated sizes (2008 estimates).

Tier#	Number	Goals	CPU	Disk	Tape
Tier0	1	Prompt Reconstruction, data custodial	7 MSi2k	2 TB	5 PB
Tier1	O(10)	Re-Reconstruction, data custodial (shared)	10 MSi2k	7 TB	10 PB
Tier2	O(50)	Simulation, analysis	13 MSi2k	5 TB	0 PB

3. Computing goals

The computing tasks in the CMS environment are numerous and demanding:

- Farming on *TierX* - which includes also operating a LCG[2]/OSG[3] Grid:
  - Verification jobs sent hourly (SAM[4]).
- Stressing CPUs, storage and geographic links:
  - JobRobot[5] to keep farms full, with fake jobs if nothing else is available;
  - LoadTest[6] to keep links operative even when nothing needs to be transferred
- Produce O(1G) simulated events per year:
  - ProdAgent[7] helps submission of O(100k) jobs a day on the GRID.
- Allow user's analisys:
  - CRAB[8] to allow user's analysis via GRID with balancing on the multiple copies of the same dataset.
- Install/maintain CMS software on sites:
  - centralized installation on the GRID worldwide.

This infrastructure has been operative (with significant development) since the data challenge performed in 2006, CSA06[9], and has been able to prove day-to-day like CMS operation, at a scale of 25% of what expected in 2008. A new challenge CSA07[10] is ongoing and will scale up to 50%. Figure 1 shows how the current rate for data movement exceeds 1 PB/month, while figure 2 shows that in the first half of 2007 CMS production teams overall produced more than 30 MEvents/month (this figure is somehow already outdated, being the capacity now more than 50 MEvents/month).

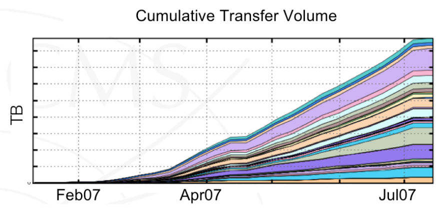

Figure 1. Data moved by CMS in the first half of 2007.

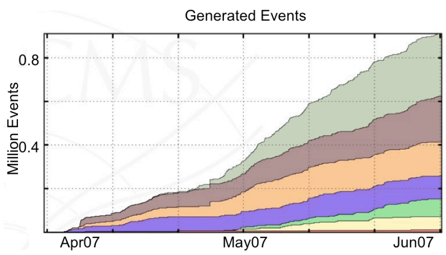

Figure 2. Simulated Events produced by CMS in the first half of 2007; the colors indicate different production teams.

Journal of Physics: Conference Series **110** (2008) 092003 doi:10.1088/1742-6596/110/9/092003

4. Offline system

The Offline group mandate is to prepare and validate the software which will be running

- on the *HLT* for the trigger,
- at the *Tier0* for prompt reconstruction,
- at the *Tier1* s for skimming and re-reconstruction,
- at the *Tier2*s for analysis and simulation.

The software development is carried out by subgroups responsible for the different aspects; among these, the most notable are *Framework*, *Reconstruction*, *Simulation*, and so on.

What is described here is the *new* CMS software, named *CMSSW*[11], which is the result of the complete rewriting started in 2005, and which supersedes the older developer's software named *ORCA*[12].

A complete description of *CMSSW* is not possible here; the most important facts are

- two simulation approaches are available: the first one based on Geant4[13], with detailed material description and well validated at test beams, the second using parametrizations and hence being 1000x faster, but still accurate enough for most tasks;
- a complete reconstruction program, featuring performance exceeding the old CMS software. The reconstruction code is suitable for both trigger and offline operation, and has been tested in data challenges and analysis efforts since more than one year;
- a visualization suite, usable also for detector commissioning;
- the skimming capability to reduce from Full Event to smaller subsets still suitable for analysis: the current AOD events are taking 50 kB/ev, close to their specifications.

Apart from all the validation plots we cannot show here, we would like to stress the enormous effort sustained by the developers to complete the transition to the new software: in less than 2 years, more than 1 million lines of code were ported/written (see figure 3), and the number of active developers has increased steadily to roughly 200 (figure 4).

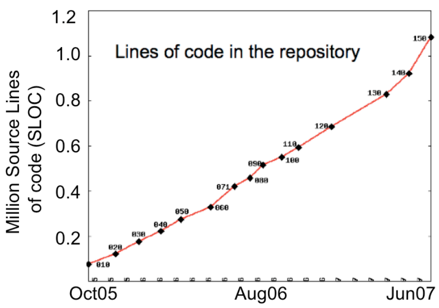

Figure 3. Lines of code per *CMSSW* release.

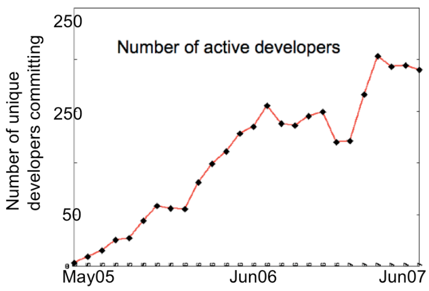

Figure 4. Number of active developers per release.

5. Conclusions

The LHC collider and CMS detector are less than one year away from data taking, and the CMS computing and offline communities are focusing on the commissioning of the relative systems. A continuous testing model for computing components, where all the components are constantly stressed, has been put in place and is working steadily since more than one year. The Offline system is increasing to a very fast rate, and when compared to 2006 one it is not only at least doubled in size, but also actively verified and checked with the previous system, and, more importantly, the available test beam data.

References

[1] http://monarc.web.cern.ch/MONARC
[2] http://lcg.web.cern.ch/LCG/
[3] http://www.opensciencegrid.org/
[4] https://twiki.cern.ch/twiki/bin/view/CMS/SAMForCMS
[5] https://twiki.cern.ch/twiki/bin/view/CMS/JobRobot
[6] https://twiki.cern.ch/twiki/bin/view/CMS/LoadTest07
[7] https://twiki.cern.ch/twiki/bin/view/CMS/ProdAgent
[8] https://twiki.cern.ch/twiki/bin/view/CMS/CRAB
[9] http://cmsdoc.cern.ch/cms/archives/07/LHCC/csa06.pdf
[10] https://twiki.cern.ch/twiki/bin/view/CMS/CSA07
[11] http://cmsdoc.cern.ch/cms/cpt/Software/html/General
[12] http://cmsdoc.cern.ch/orca
[13] Nuclear Instruments and Methods in Physics Research A 506 (2003) 250-303.