

Experience with CMS Offline and Computing in LHC Runs 2010-2011

Jean-Roch Vlimant¹

University of California Santa-Barbara

Santa Barbara, CA 93106, USA

E-mail: vlimant@cern.ch

In this presentation we will discuss the experience with the CMS computing model during the LHC runs 2010-2011. We will outline how the experiment operations has evolved during the first few months of operations. The current state of the Offline and Computing projects will be presented and we will describe the initial experience with active analysis users and real data. We will include Tier-0 processing, reprocessing steps on data and fast turn-around calibrations. We will address the issues that worked well in addition to identifying areas where future development and refinement are needed.

*The 2011 Europhysics Conference on High Energy Physics-HEP2011
Grenoble, Rhône-Alpes*

July 21-27 2011

¹ Speaker

1. Introduction

The CMS Computing model has a tree-like structure, with tiers at each generation. The root being Tier-0, at CERN, close to the CMS experiment. There are 7 Tier-1 sites, across the world (see Fig. 1). There are currently 52 Tier-2, even more widespread, and a numerous amount of Tier-3. While Tier-0 and Tier-1 are dedicated to production, Tier-2 and Tier-3 are shared resources shared with CMS users.



Figure 1: Location of the CMS Tiers of the computing model.

In this proceedings, we will detail the CMS Data production and distribution model, then report on the CMS software management, quality monitoring, and eventually the availability of dataset for analysis during the years 2010-11.

2. CMS Data Production and Distribution

2.1 Detector Data

The RAW data coming from the detector is handled on the close site (Fig. 2), and transmitted to the Tier-0 for repacking, express processing and prompt reconstruction. RAW dataset are archived at CERN and duplicated for the purpose of data reprocessing and backup to custodial Tier-1. Prompt calibration and quality monitoring of the CMS data is performed from the express reconstruction, performed at Tier-0. Prompt reconstruction is performed at Tier-0 with a delay, to accommodate for prompt calibration. Certification (masking of bad data) is derived from the prompt reconstruction. Dataset are send in both detailed (RECO) and analysis level (AOD) descriptions to the Tier-1 sites. Data reprocessing is performed at Tier-1 from the custodial copy of the RAW data. Skimming of the data, for Physics analysis purpose, is performed at Tier-1.

2.2 Simulated Data

The event generation for simulation, originally designed to be performed solely at Tier-2 (Fig. 2) has been performed at both Tier-1 and Tier-2 for a better utilization of resources. The archiving of the generated event dataset is done at Tier-1. The rest of simulation and reconstruction is performed either at Tier-1 or Tier-2 depending on the type of production campaign (full simulation or reprocessing).




Figure 2: Flow of data within CMS Tiers. Data in red, Simulation in green.




Figure 3: BreakDown of CMS Data traffic in PheDEx in 2010

2.3 Utilization

Regardless of the dataset, RECO datasets are archived at Tier-1 (Fig. 2), while analysis level dataset are distributed to Tier-2. All analysis jobs are ran on Tier-2 and Tier-3. The data transfer system (PheDEx) has reached a steady rate of a couple of 100 of terabytes per day, over the last few years, and dataset availability to users has been greatly improved by commissioning Tier-2 to Tier-2 transfer links (Fig. 3)

3. Software Management

The CMS software for simulation and data processing follows a standard integration mechanism, with daily builds, performance monitoring,... Software integration closure test and physics performance monitoring is performed frequently, forming a good base for software development. There is a strong core of software experts managing the framework for modular contribution to an ever complex simulation and reconstruction software. Development and understanding of CMS data is facilitated by the CMS event display, product of the consolidation of two previously distinct projects, available under several form depending on the targeted use.

The monitoring of the quality of the CMS data is performed at three levels of the data taking, utilizing shifters and user friendly interfaces for data certification and detailed investigation. This forms a strong platform for software validation.

Major release versions of the CMS software are targeted for production of consistent simulation and data datasets for Physics conferences. While data reprocessing during in 2010-2011 has been a matter of weeks, simulation reprocessing requires months and therefore needs careful scheduling to ensure availability. Integration of evermore precise simulation, geometry, event generator, reconstruction algorithms and calibration need to be scheduled accordingly.

4. Dataset Availability

Five major release have been deployed at Tier-0 in 2010-11 to process about 2 billion detector events, with a very low failure rate ($\sim 10^{-8}$). While RECO datasets were distributed centrally in 2010 for commissioning, AOD has been distributed in 2011 for analysis, in order to limit the amount of storage required at Tier-2. The utilization of the Tier-0 farm varies with the

instantaneous luminosity delivered by the LHC. Full occupancy for a few days has been reached with late spring 2011 fills. Further optimization on both the resource management and the deployed application are ongoing in order to cope with the increase in luminosity.

Up to 4 major releases and 2 calibrations were deployed for data reprocessing in 2010, demonstrating the ability of the CMS computing system to deliver updated dataset as fast as possible. In the year 2011, up to this conference, only one release had been deployed.

Major campaigns (Fig. 4) of event generation generally happen once per year. Other campaigns consisted in updating the detector simulation, modify the pile-up profiles and follow the significant changes in event reconstruction software. Throughout production, consistent datasets were provide to the collaboration for prompt analysis. The availability of samples for analysis has been made possible by a careful and educated scheduling of production.

Both detector data and simulated events were accessed by the user via the GRID. Analysis update were made possible with constant effort from the CMS Offline & Computing project to provide and improve tools to use distributed resource. Analysis oriented releases were created as evolution of production releases to incorporate end-user calibrations, analysis tools, ... so as to create a user friendly environment for analysis. Tier-2 sites were shared between users and Monte-Carlo production, with almost 80% sites ready for computing and an increased user share thanks to shifting of production to Tier-1, more and more resource are made available for Physics analysis.

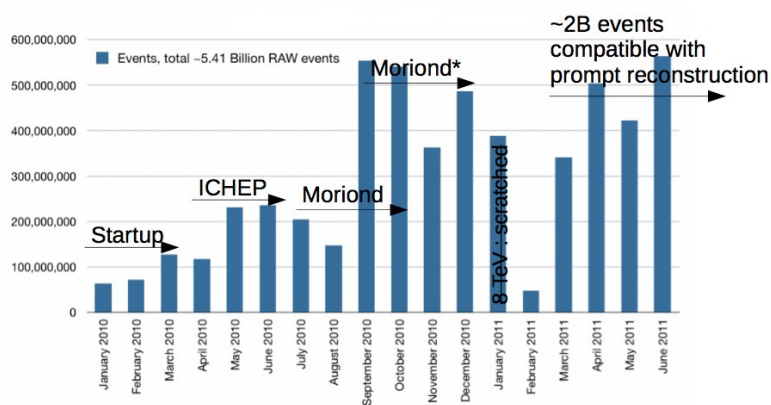


Figure 4: Monte Carlo Production in 2010-11 [Simulated Events per Month]

5. Conclusions

We report in this proceeding on the experience with Offline and Computing projects in CMS for the years 2010-11. The data from the CMS detector, so as to simulated data have been handled efficiently and delivered to the CMS collaboration for Physics analysis. The Offline project has been able to manage the necessary software for event simulation, data processing, shared resource usage, physics visualization. The mode of operation has been constantly adapted to the LHC agenda and excellent performance, and is transitioning, with upgrades, to stable data taking operation with high pile-up at the LHC.