CMS Event Display and Data Quality Monitoring at LHC Start-up

Ianna Osborne[1,*], George Alverson[1], Giulio Eulisse[1], Shahzad Muzaffar[1], Lucas Taylor[1], Lassi Tuura[1], Nicola Amapane[2,3], Riccardo Bellan[2,3], Gianluca Cerminara[3], Domenico Giordano[4], Norbert Neumeister[5], Chang Liu[5], Gianni Masetti[6], Maria-Santa Mennea[7], Giuseppe Zito[7], and Ilaria Segoni[8]

[1] Northeastern University, Boston, MA, USA
[2] Università di Torino
[3] INFN, Torino, Italy
[4] Universita' & INFN di Bari, Italy
[5] Purdue University, West Lafayette, USA
[6] Universita' di Bologna e Sezione dell'INFN, Bologna, Italy
[7] INFN Bari, Italy
[8] CERN

[*] To whom any correspondence should be addressed. E-mail: Ianna.Osborne@cern.ch

Abstract. The event display and data quality monitoring visualisation systems are especially crucial for commissioning CMS in the imminent CMS physics run at the LHC. They have already proved invaluable for the CMS magnet test and cosmic challenge. We describe how these systems are used to navigate and filter the immense amounts of complex event data from the CMS detector and prepare clear and flexible views of the salient features to the shift crews and offline users. These allow shift staff and experts to navigate from a top-level general view to very specific monitoring elements in real time to help validate data quality and ascertain causes of problems. We describe how events may be accessed in the higher level trigger filter farm, at the CERN Tier-0 centre, and in offsite centres to help ensure good data quality at all points in the data processing workflow. Emphasis has been placed on deployment issues in order to ensure that experts and general users may use the visualization systems at CERN, in remote operations and monitoring centres offsite, and from their own desktops.

1. Introduction

Heading toward the LHC startup the CMS offline group is preparing and testing its final software systems especially the tools for the data quality monitoring. An important component in monitoring the quality of data in any experiment is an event display. This software provides a visual interpretation of the data and detector on an event-by-event basis. The deployment and operation of the CMS event display has started for the CMS magnet test and cosmic challenge (MTCC) [1]. The ongoing sub-detector commissioning and the monthly global runs give an invaluable feedback to the developers. This operation in a nearly running experiment experience helps to build a stable and robust system, which should operate with a minimum maintenance effort for years.

The CMS event display is implemented based on the “Interactive Graphics for User ANAlysis” (IGUANA) [4, 5, 6] visualization framework and toolkit and is fully integrated with the CMS software framework.

2. CMS Display Operation during MTCC

During the MTCC a 20 degree combined sub-detectors slice of CMS with a magnet has triggered and recorded the cosmic rays, checked the noise and the inter-operability using as near as possible final readout and the auxiliary systems. The 24/7 CMS operation procedures have been tried out which encouraged the use of the final systems as far as possible, deliberately limiting any MTCC specific development. Around 25 million “good” events have been recorded with at least the DT triggers and ECAL + TK in readout, including a 15M events at a stable field ≥ 3.8T. The data-taking efficiency reached over 90% for extended periods.

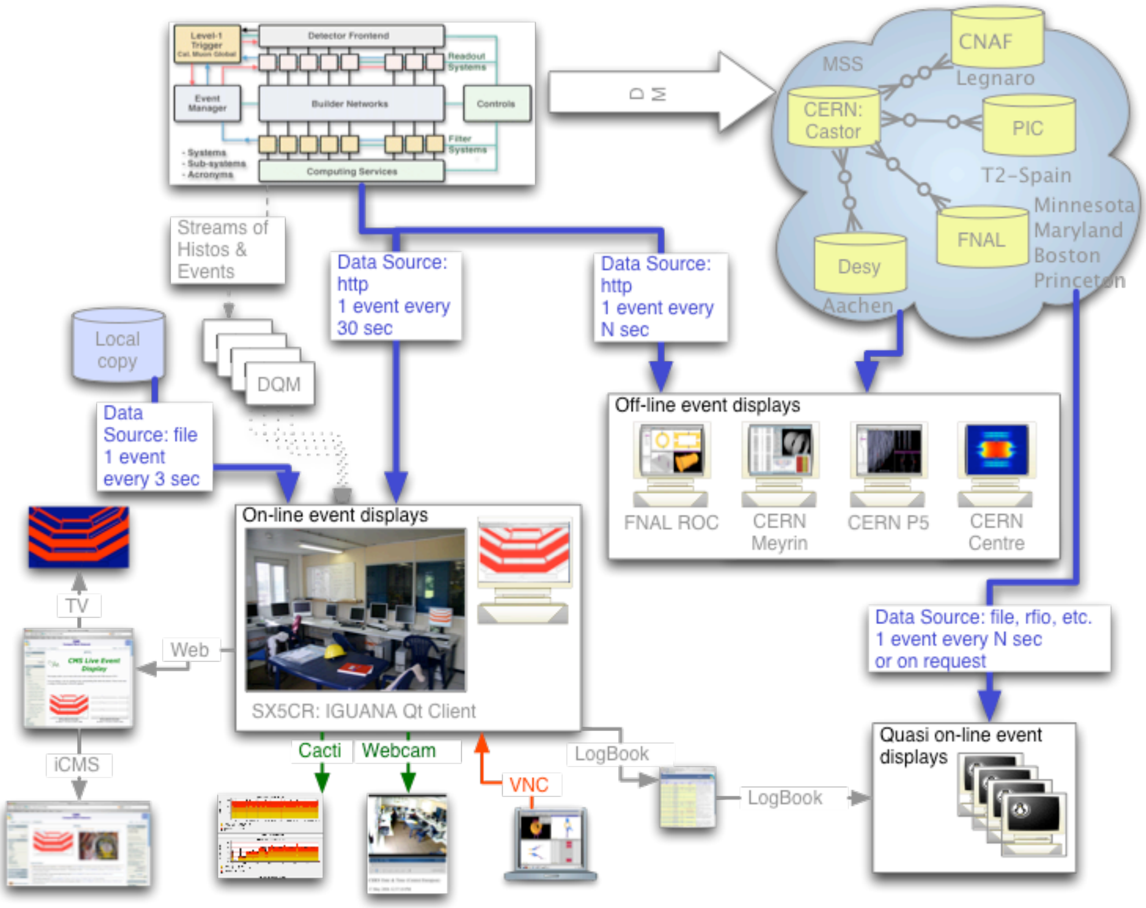

Figure 1: The CMS event display workflow during MTCC.

The real time feedback from the CMS event display and the Data Quality Monitoring system helped in tuning and timing-in the sub-detectors. Both systems were deployed at the LHC Point 5 (P5) control room (see Fig.1) for the 27/7 MTCC operations.

Multiple instances of the running event displays can be divided into three categories:

- On-line event displays which connect directly to the central DAQ data streams and display data in real time.
- Quasi-on-line event displays receiving the data from the local files within local network with a very short latency.

- Off-line event displays which get the data from the Tier 0 or Tier 1.

An on-line event display is installed in the control room on a private network. The connection to the data source is implemented via an http connection to the storage manager running in the central DAQ. Whenever the storage manager was not running the event display would fetch the data from the local files or an events playback server. The on-line event displays regularly produce the snapshots and publish them to a Web server.

The MTCC quasi on-line event displays, such as the custom tracker event display, the event displays at ROC, etc. are installed on a public network and may experience certain latency in accessing the latest data.

The off-line event displays access the data distributed by the Data Management systems.

The event display operation at P5 is monitored via Cacti and a web-cam. It is remotely controlled via VNC. The event display operator on shift responsibility is to submit regular reports to a LogBook.

3. What is IGUANA

IGUANA is a visualization environment. It allows a user to build his or her visualization application at run-time, for example, CMS event display(s). IGUANA is based on plug-ins a.k.a. iglets.

An iglet is a software module that adds a specific feature(s) or service(s) to an IGUANA application. An iglet-based event display is fully configurable at the start-up, and its functionality is defined by the capabilities of the loaded iglets (See Fig. 2).

A text label assigned to each iglet at compile time can be used in a configuration file to load the iglet at runtime.

4. IGUANA and CMSSW

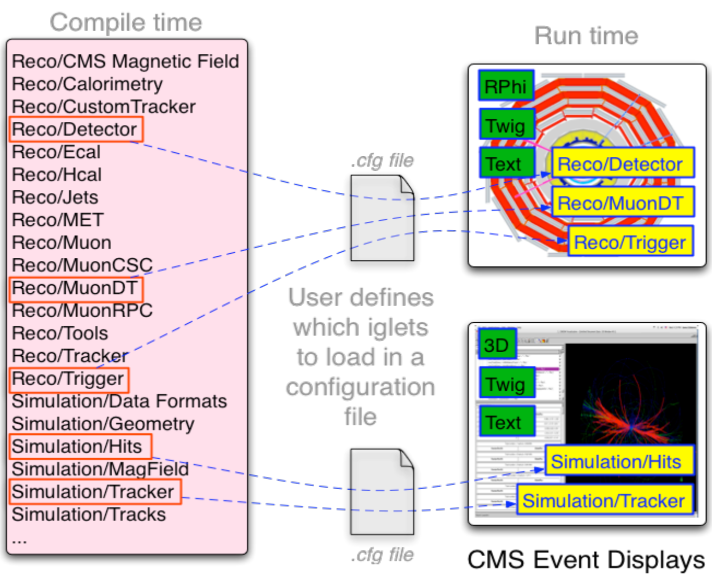

Figure 2: CMS Event Display loads listed in the configuration file iglets to define its functionality.

4.1. CMSSW Event Data Model

The CMSSW event data model [5] requires that all event data processing modules communicate only through a single data structure called the "Event". The modules are executed according to the schedule specified in the job configuration path. A special input source module can be used to insert in the Event the raw data received from the CMS DAQ. Thus, the offline software can run in a completely transparent fashion both in purely offline applications and in the HLT.

The reconstruction modules get from the Event the needed input data collection and produce the higher level reconstructed data collections which become part of the Event. The Filter modules are used online to avoid forwarding events to the Storage Manager application, whose task is to collect accepted events from several FUs nodes and assign them to event streams identified on the basis of the HLT pattern.

4.2. Event Data Collection Display

The event display is fully integrated with the CMSSW framework: it instantiates and runs the Event Processor - just as a usual CMS analysis job does. Just as the analysis job, the event display needs a configuration file to define a process with a data source and/or Event Setup. The event display uses the CMS Event Data Model to display Event data and allows the user to add CMSSW modules to the process path to produce and display new transient data collections. The event display takes the auxiliary information from Event Setup (See Fig. 3).

The configuration file can also define the event display configuration, for example, a list of iglets to load.

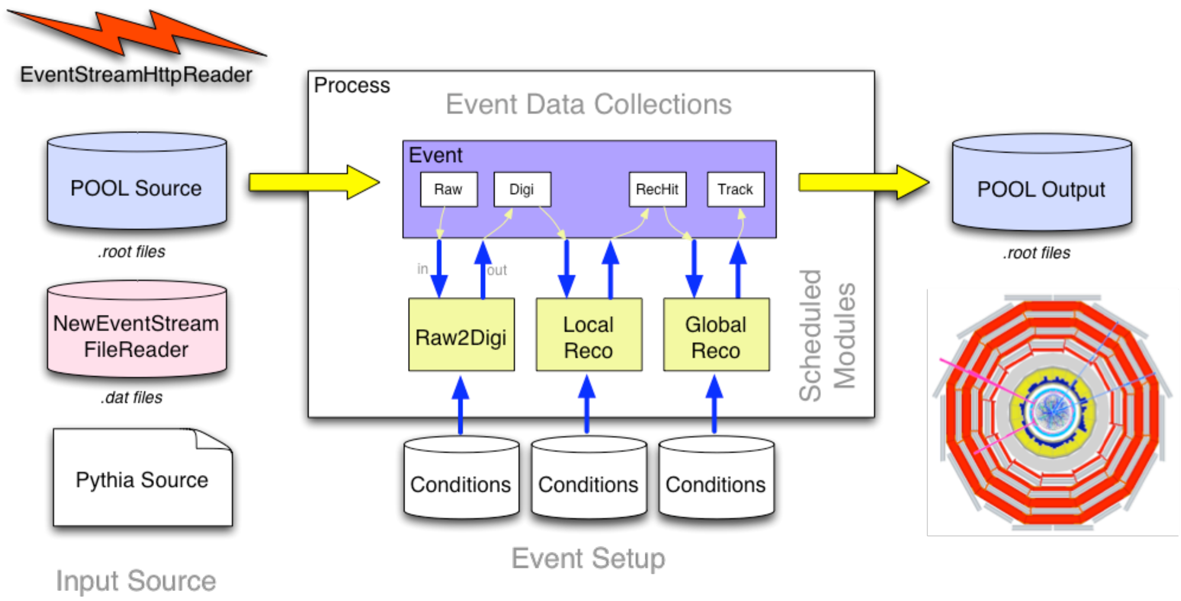

Figure 3: **CMS Event Processing**

IGUANA runs on a configuration file: a process, which defines the data source, the Event Setup producers and the modules schedule. The data source or the data file has to be accessible from the system where the event display is run from - preferably within a very short latency. The event display dynamically discovers the Event content, it retrieves all event data collections from the data source.

The data collections within the Event are uniquely identified by four quantities:

- C++ class type of the data
- A module label
- A product instance label
- A process name

The event display retrieves a complete list of the data collections from the data source and presents them in a hierarchical tree. The data collection type is used to register an iglet capability. A Twig, a C++ class, represents every data collection.

There can be more then one Twig for the collection data type defined in the different iglets.

There can be more then one collection represented by a single Twig. Every Twig defines how it displays the data collection within its Twig::update member functions.

At startup a user decides which iglet to load to get a specific representation (See Fig. 4).

If a collection cannot be displayed, it means that an iglet capable of displaying this particular data type has not been loaded.

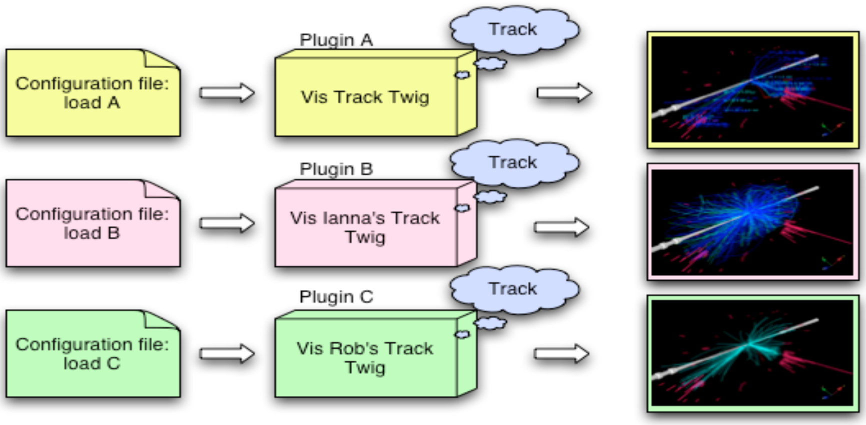

Figure 4: Dynamic Track representation.

4.3. Event Setup Display

The Event data constituencies usually belong to a detector unit and positioned in its local coordinate system. IGUANA displays the data and geometry in the global coordinates. Local to global transformation is done via Event Setup (See Fig. 5).

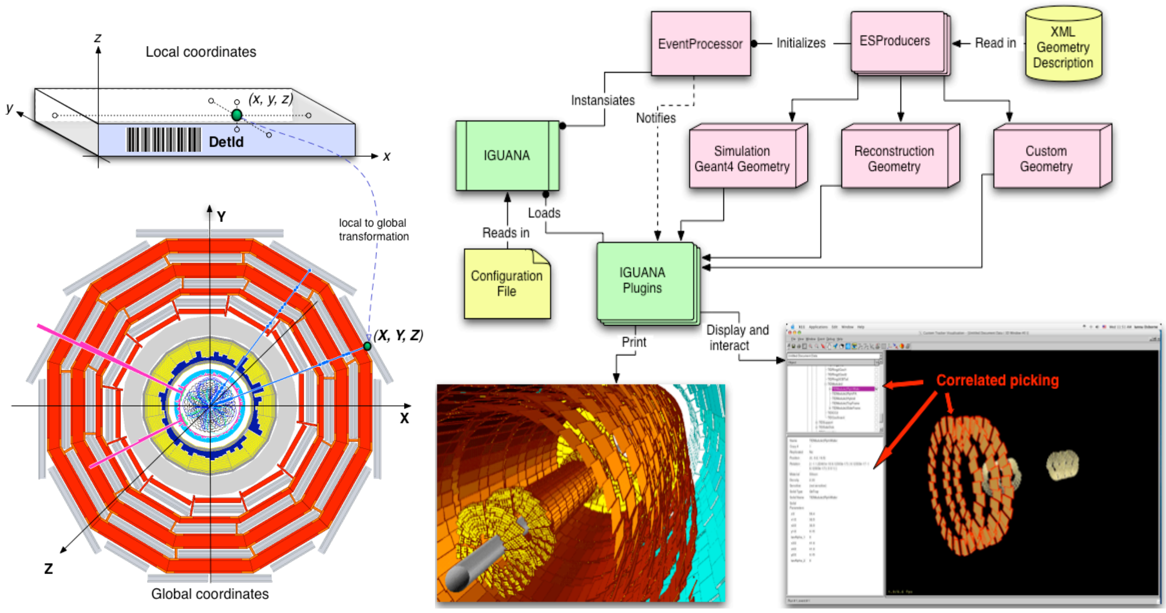

Figure 5: Local to global coordinate conversion is done via Event Setup. The iglets are notified when the Event Setup producer delivers the geometry.

The configuration file defines which Event Setup producers deliver the geometry for the application.

It is possible to have more then one geometries producer in the same application. For example, if it is necessary to compare the simulated geometry with the reconstructed one. The definition of the geometry comes from an XML description. The same event setup producer delivers either the MTCC geometry or the reconstructed geometry. That is why these geometries cannot be used together in the same application.

5. Statistical Data Quality Monitoring

Complimentary to the event-by-event view of the detector the cumulative information gives a very good overview of what's going on.

6. Summary and Outlook

How to present the amount of information in a most intuitive way? Learn from previous experiments, learn from the detector commissioning, and learn from the global runs. Use multiple and specific views developed by the detector people, involve physicists in the development process as early as possible. Deploy tools early.

Is it necessary to display all the channels? Never say never. Keep room for expanding requests.

Does data quality monitoring on event-by-event basis guarantee a quick turn around? IGUANA itself does not guarantee good quality data, but it does guarantee an honest view of the detector response in real time.

References

[1] CMS Collaboration 2006 The CMS Magnet Test and Cosmic Challenge (MTCC Phase I and II) Operational Experience and Lessons Learnt, CMS NOTE-2007/005

[2] G. Bruno, Software for the CMS Cosmic Challenge, CMS CR 2006/017

[3] CMS Collaboration 2002 The Trigger and Data Acquisition project, Technical Design Report, Volume 2: Data Acquisition & High-Level Trigger CERN/LHCC 2002-26

[4] CMS Collaboration 2006 CMS Physics Technical Design Report, Volume 1, Section 2.10: Visualization

[5] V. Innocente, G. Eulisse, S. Muzaffar, I. Osborne, L.A. Tuura, L. Taylor, Composite Framework for CMS Applications CHEP04, Interlaken, Switzerland, September 27–October 1, 2004.

[6] G. Alverson, G. Eulisse, S. Muzaffar, I. Osborne, L.A. Tuura, L. Taylor, IGUANA Architecture, Framework and Toolkit for Interactive Graphics CHEP03, La Jolla, California, March 24–28, 2003.

[7] M.S. Mennea, I. Osborne, A. Regano and G. Zito, CMS tracker visualization tools, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment Volume 548, Issue 3, 21 August 2005, Pages 391-400

[8] CMS Visualisation Tools, M. S. Mennea, A. Regano, G. Zito, I. Osborne, CHEP04, Interlaken, Switzerland, September 26-October1, 2004.