

The Web Based Monitoring Project at the CMS Experiment

**Juan Antonio Lopez-Perez¹, William Badgett¹, Ulf Behrens², Irakli Chakaberia³,
Youngkwon Jo⁴, Kaori Maeshima, Sho Maruyama¹, James Patrick¹, Valdas
Rapsevicius¹, Aron Soha¹, Mantas Stankevicius¹, Balys Sulmanas¹, Sachiko Toda³
and Zongru Wan³**

¹Fermi National Accelerator Laboratory, Batavia, IL 60560, USA

²DESY, Hamburg, Germany

³Kansas State University, Manhattan, KS 66506, USA

⁴Korea University, Seoul, KR

E-mail: patrick@fnal.gov

Abstract. The Compact Muon Solenoid is a large a complex general purpose experiment at the CERN Large Hadron Collider (LHC), built and maintained by many collaborators from around the world. Efficient operation of the detector requires widespread and timely access to a broad range of monitoring and status information. To the end the Web Based Monitoring (WBM) system was developed to present data to users located anywhere from many underlying heterogeneous sources, from real time messaging systems to relational databases. This system provides the power to combine and correlate data in both graphical and tabular formats of interest to the experimenters, including data such as beam conditions, luminosity, trigger rates, detector conditions, and many others, allowing for flexibility on the user's side. This paper describes the WBM system architecture and describes how the system has been used from the beginning of data taking until now (Run1 and Run 2).

1. Introduction

The Compact Muon Solenoid (CMS) [1] is a general purpose detector at the Large Hadron Collider (LHC) [2] at CERN built to study proton-proton, proton-heavy ion, and ion-ion collisions at center of mass energies up to 14 TeV. The operation of CMS involves tens of millions of sensor channels, substantial supporting infrastructure, and complex trigger and data acquisition (DAQ) systems. The vast and diverse information to monitor includes detector and environmental conditions, DAQ status, run configuration, trigger rates, luminosity, and accelerator parameters. Extensive monitoring is important for efficient and high quality data taking. While there exist dedicated monitoring applications for personnel at the experiment site these applications are generally subsystem specific and not easily accessed remotely. The CMS collaboration is truly global, with over 3000 collaborators from over 175 institutions from over 40 countries. With such distributed expertise and a variety of data sources, there is a need for unified remote access to monitoring data in order to identify and correct issues with data taking in a timely manner.

The goal of the CMS Web Based Monitoring (WBM) system is to provide collaborators a way to monitor the operational status of the detector and diagnose problems from any location via a web browser. The WBM is a set of hardware and software that acquires data from diverse online monitoring sources and makes them available to a set of web based applications accessible to any authenticated user

anywhere. These applications give access to a summary of the current status of the experiment as well as convenient access to historical data. A single application can easily include information from different major online systems and the LHC accelerator. These applications are used by shift crew members, detector subsystem experts, operations coordinators, and those performing physics analyses. The tools have become important in managing data taking operations for the CMS detector. The monitoring described in this paper is complementary to the CMS Data Quality Monitoring (DQM) tools that are based on event data [3].

This paper will describe the WBM infrastructure in section 2, core applications providing overall summary information and also some more specialized tools, and some of the detector subsystem monitoring applications in section 3, and experience with the system during data taking and future plans in section 4. Section 5 contains a summary and conclusions.

2. System Architecture

The primary systems associated with operation of the CMS detector are the trigger and data acquisition systems [4], the detector control system (DCS) [5, 6], and the LHC accelerator. These systems are built on different software frameworks. The goal of the WBM system is to provide applications that may include data from any or all of these sources. To accomplish this goal, the system is based on a 3-tier architecture, a simplified view is shown in figure 1. Data flows from left to right in this picture. The servers, *Scaler* and *Lhcgmt*, acquire data from these sources and write it to a database and make some of it directly available in real time. A web based application architecture is used to access and display the data. This has the advantage that applications can be accessed from anywhere via a standard web browser without installing any custom software. Also applications do not need to know the specifics of how the data are acquired, they obtain it from the database or via basic real time subscriptions as described below. And as the applications do not directly access the detector systems there is little chance of interference with operations. The local online web servers provide access to users on the local experiment network. The applications are accessible outside of the experiment network via a proxy server on the CERN general network. To reduce the load on the online web servers, dedicated offline WBM servers operate on a copy of the online database and on data from the CMS offline database. These services are accessible for collaborators worldwide. The WBM system is compliant with CERN security standards.

2.1. Data Acquisition

Accelerator information is acquired by the *Lhcgmt* system via a hardware interface to the LHC General Machine Timing (GMT) system, [7] and also via the Data Interchange Protocol (DIP) [8, 9]. The GMT system provides synchronous event timing and basic status information with high reliability over dedicated serial links. The status information includes the LHC machine state, beam energy, and intensities. The information is logged to the database and made available to the DCS and data acquisition systems. The beam status is used for example to record periods of stable beams, and to put the detector in a safe state when beam conditions are not stable. More extensive accelerator information including bunch by bunch measurements is acquired via DIP from the LHC control system. DIP is a publish/subscribe protocol that allows data to be exchanged between loosely coupled systems. Also luminosity and beam losses measured by the CMS detector are published to the LHC DIP network where they are available to the machine control system.

The *Scaler* server deals with data from the DCS, and the trigger and event data acquisition systems. Information such as high voltage and environmental status of detector subsystems is published via DIP by the DCS. Status of the trigger and DAQ systems is provided via the monitoring system of the online data acquisition framework XDAQ [10]. XDAQ is a software platform designed at CMS specifically

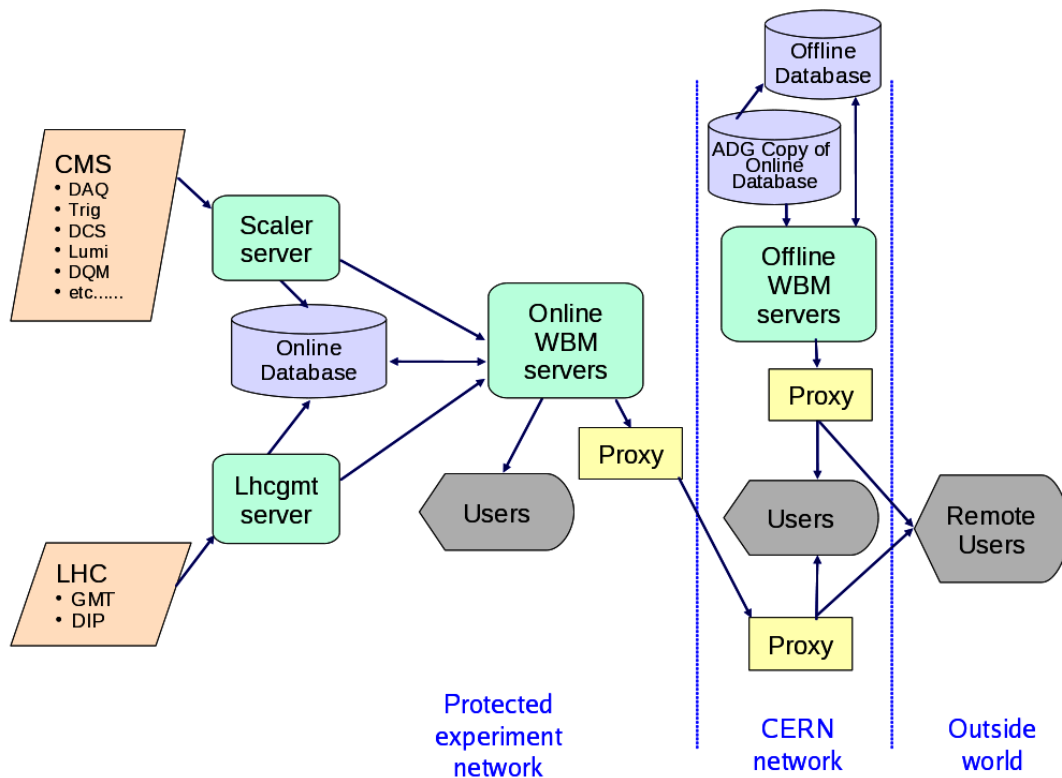


Figure 1. Overview of the dataflow in the CMS WBM system

for the development of distributed data acquisition systems. In addition to obtaining general status information, the *Scaler* server examines rate-counter information to note starts and stops in data taking and writes this information to the database. This information is used by WBM services that account for operational down time. Trigger scaler data is also recorded in the database. The *Scaler* server is a front-end in the data acquisition system, and contains hardware interfaces allowing it to insert some data it acquires into the raw data stream. This data includes the beam position, luminosity, bunch number, scalars, and subsystem status and other information required by the High Level Trigger and DQM systems. Online luminosity information measured by CMS subdetectors is obtained from the luminosity subsystem via DIP and also logged to the database. This information is used to mark blocks of delivered luminosity, approximately 23 seconds in duration, called Luminosity Sections. The detector subsystem status as well as scalars and other run time information is recorded in the database for each of these sections. The Luminosity Sections where some part of the detector is not functioning properly may be excluded from some offline physics analyses. These sections are long enough to allow a precise determination of the luminosity but short enough so that data loss is minimized when the detector is not fully functional.

Note that many of these functions performed by the WBM data acquisition infrastructure are integral to the routine recording of CMS event data in addition to the monitoring data for WBM applications.

2.2. Databases

The CMS online database contains a wealth of information regarding the current and past status of the CMS detector written by a variety of sources in addition to the information logged by the WBM system [11]. It is an Oracle database, with several servers working within the protected CMS experiment network. The stored data is copied to a database server on the general network via Oracle Advanced Data Guard [12] so that it is available to the offline WBM servers. The WBM server configuration as implemented provides globally useful, but carefully controlled, access to the data for people at the CMS site, as well as for collaborators at remote locations.

2.3. WBM Computing Infrastructure

The WBM infrastructure includes dedicated machines for the data acquisition functionalities described in section 2.1. In addition, to serve two distinct user bases, three web servers are used. Two servers reside within the CMS protected network for online users, and one server on the CERN general network for other users. For web service development, testing, and as an emergency backup, two additional servers configured in the same way as the production servers are used within each of the CMS and CERN networks. All computers run the 64-bit version of Scientific Linux CERN (SLC) [13].

As the WBM system is important for CMS operations, a number of steps are taken to ensure its reliability. The local computing infrastructure is monitored as part of the general CMS online system. In addition, the WBM system employs scripts to check the Tomcat log files for errors and perform a restart if needed. These scripts also look for additional WBM specific issues and notify experts if problems are found. In case of failure of one of the main servers, one of the backup systems can be readily substituted. The databases and offline WBM computers are monitored by the CERN IT department and periodic reports are issued.

2.4. Software Implementation

The WBM web server runs an Apache daemon to respond to HTTP requests. Most of the user interactive work is handled through Apache/Tomcat [14] using a basic Java servlet architecture. These packages are open-source and widely used in the scientific community. They provide sufficient infrastructure for WBM so that development focuses on the interactive web services. The main Apache daemon serves a small number of static pages and forwards requests to the Tomcat engine. JDBC and SQL are used to read data from the database. Some servlets subscribe to real time DIP data when an application with auto-refreshing functionality requests current values. While much of the code is written in Java, some data analysis and production of static plots is performed by the C++ based ROOT package [15] with output available as PNG images or ROOT files that can be saved for further analysis (*e.g.*, trigger cross section dependence as a function of instantaneous luminosity). The JFreeChart [16] library is used to produce some PNG plots within the Java language framework when the other data formats are not required.

On the client side, data are presented in web pages using basic HTML, CSS, and JavaScript that is supported by most modern web browsers. Development and testing concentrates on the recent versions of the Firefox browser on Linux and MacOS most often used to access WBM applications. Issues found with other browsers are addressed as needed. An example is CSS rendering differences in the Safari browser on MacOS. AJAX and XML are used to do automatic refreshing of real-time displays. Some client services are in the form of Java applets to provide interactive features to users. In the latter part of the initial LHC run, some of these applets were migrated to HTML5, jQuery [17], and Highcharts [18] which avoids various security and other issues with Java applets. These packages provide interactive charts, allowing users to change axis ranges, zoom in and out, and mouse over data points to see the values.

The WBM server provides an API so that some of the stored data may be easily fetched by non-WBM applications. The data are provided in XML format in response to an HTTP query containing the desired items and date or time range.

3. WBM Services

The WBM services are accessible through a simple web page that provides a single point of entry. There are monitoring services that focus on the both real-time and historical status of the detector and accelerator. Other tools track data taking efficiency, and perform other functions helpful for detector operations. Most services are linked to other services to allow users to navigate quickly to other relevant information with few mouse clicks. In what follows a partial set of the applications in the WBM system is summarized in four categories: real-time monitor, summary and report, sub-system specific applications, and generic services, for the corresponding primary users: online and offline shift crews at the experiment site and remote sites, operation managers, online, and offline sub-system experts, and generic users, respectively.

3.1. Real-time Monitors

The WBM real time monitor services include a basic one screen non-interactive overview of the detector and accelerator status. The *Page Zero* and *CMS Page 1* pages concisely present the most important operational information: the LHC energy and luminosity, trigger rates, the CMS status of the data taking run, and a summary of detector subsystem statuses. The latter page is publicly accessible to anyone via the CERN Vistar page [19].

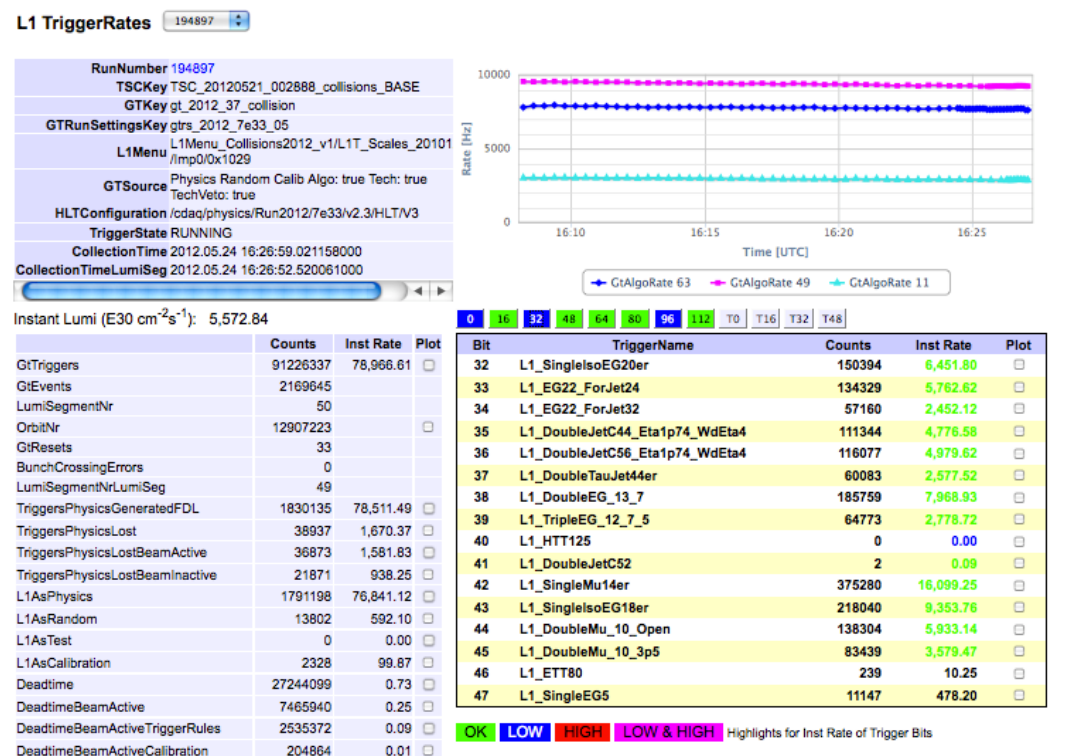


Figure 2. The *Trigger Rates* WBM application. This shows real time or historical information on overall trigger rates and dead times, as well as individual trigger rates. Any of this information can be selected for plotting.

More detailed real time information is displayed in dedicated applications such as the *Lhc Monitor* and *Trigger Rates*. The *Lhc Monitor* service presents the LHC parameters: the most recent status, a list of the most recent GMT events, and basic parameters for each beam, obtained from the GMT system. The *Trigger Rates* service displays trigger counts, rates, and configurations in real time in both tables and plots as well as historical data from the database for comparison (see figure 2).

3.2. Summary and Report Services

Available accelerator and data taking information are systematically organized and presented in terms of arbitrary data-taking periods, LHC fills, and CMS runs with the corresponding summary and reporting applications: *Data Summary*, *Fill Report*, and *Run Summary*. These applications provide a concise summary of the detector performance in different time scales of interest. For example, the *Data Summary* service displays brief summary information and plots of instantaneous and integrated luminosity and data taking efficiency by a day, week, year, or user specified range. The *Fill Report* and *Run Summary* applications, linked from the *Data Summary*, contain more detailed information about the LHC fills and CMS runs, such as tables and plots of luminosity, backgrounds, accelerator information, and trigger rates as well as links to more information about them.

For the evaluation of data taking efficiency, dedicated applications are provided. The *CMS Run Time Logger* (RTL) reports real time and historical operational efficiency details in tabular and graphic forms using information on gaps in data taking logged by the *Scaler* server described in section 2. The reasons for down times are recorded by the shift crew members using a standalone application at the experiment site.

3.3. Sub-system Specific Applications

Operating parameters (*e.g.*, voltages and currents) and environmental conditions (*e.g.*, temperatures and humidity levels) of sub-detectors are presented in a hierarchical tree or by pictorial detector geometry by the *Last Value* application, with an option to plot these variables as a function of time. Each sub-system group can define their own applications correlating sub-system specific information and the data taking status of the detector. These services are created by subsystem personnel by taking advantage of the WBM servlet and data acquisition architecture. It is straightforward to add new data items to be acquired, and to add new servlets and displays that have access to all the other status information in the database. Some subdetector groups have developed an additional local infrastructure that minimizes the programming required to add new plots or other summary pages. For example, the Tracker subsystem uses template programs based on ROOT to query the database and create plots. The Hadron Calorimeter and Pixel subsystems use “portlets” [20] that provide a modular infrastructure for their displays. The breadth of information available to the WBM services makes possible a higher-level analysis of the data. For example, the *TXMon fitter* service plots trigger cross sections as a function of luminosity. This service allows fits to be performed and the results used to predict and monitor rates for future runs, and also to design future trigger configurations for the Trigger group.

3.4. Generic Services

Given the many diverse data sources collected at CMS, there could be subtle correlations among a set of recorded quantities some of which may not be obvious. The *Condition Browser* application allows users to visualize any of the more than 4000 quantities that are stored in the database in the form of values and time stamps. Trends over time can be plotted for any given time range, where the time stamps may be actual times, or a range of run numbers or luminosity sections. Correlation plots between any two variables can be created by joining the values with the shortest time difference.

There are some applications that take advantage of the WBM infrastructure even though they are outside of the core online monitoring scope of the system. The *CMS Run Registry* (RR) and DCS alarm systems are examples of such applications.

The RR is a suite of tools for tracking the quality of data taking runs for the experiment. WBM along with event quality monitoring information is used by the DQM group to determine which data taking runs or parts thereof are used by various physics analyses. The RR consists of a suite of applications deployed in the WBM servers. These include online and offline user interfaces and tools for DQM shift personnel to make the quality assessment. The general RR application provides a user interface and API for end users to browse, query, and export certification results for use in physics data analysis or other uses. The RR is also used to trigger the calibration sequence based on data acquired during a run. When a new run is taken, WBM run summary information is propagated to the Run Registry. RR determines the type of the run, such as "collision run" or "cosmic run" using the information provided by the WBM system. Through RRAPI, this information (new run, type of run) is used to trigger the CMS calibration sequence of sub-systems automatically.

The WBM system is used to combine DCS status information with the accelerator state to generate a central alarm in the control room in case of high voltage or other problems in some subsystem. When an alarm is raised in the DCS system while Stable Beams are declared, a set of WBM scripts will raise an on-screen warning message and sound a distinctive alarm in the control room.

4. Experience in LHC Operation

The WBM system has proven to be valuable during CMS data taking to date as evidenced by the widespread use of summary tables and plots offline, and real time monitoring services online. During this running period there were an average of approximately 2600 unique visitors and 650,000 page views per month. The WBM architecture was able to handle the required data throughput, and down time of the system was minimal. Information from the WBM system was used by international collaboration members providing support for detector operations, and was an important tool for the run coordination team to effectively monitor and improve data taking performance. In particular, the RTL was useful in identifying sources of inefficiency in collecting data. Significant effort was then spent in correcting problems causing these inefficiencies. New features were continually added in response to operational needs. As the luminosity steadily increased during the run, the trigger rate fits were valuable in planning the trigger menu so as to minimize detector dead time. In addition to the core team, a broad range of developers were able to contribute to the subsystem portions thanks to expandability of the WBM architecture.

Following this data taking run improvements were made to the LHC accelerator to allow it to run at an energy of 6.5 TeV per beam. Operation resumed in 2015. Many upgrades were also made to the CMS detector during this time. These included major changes in the trigger, timing and data acquisition systems [21] [22] as well as to the luminosity measurement system and other detector subsystems. Many of these required significant changes in the WBM software. This was done in such a way as to minimize changes in the user interface, and to preserve the ability to view information from the initial LHC operational period. More use was made of HTML5 and the Highcharts plotting package to eliminate problems associated with Java applets mentioned in section 2. A more general API to access data stored by WBM for either external use or to simplify creation of new applications was generated. For example, it would be possible to include information from the DQM system in WBM plots via API calls. The infrastructure was modified to better implement a REST style architecture. Additional features requested by the run coordination team and the subsystems were implemented. And general updates of computer hardware, operating systems, and software packages to more recent versions were performed.

5. Conclusions

To meet the critical need for monitoring of detector status and performance, the Web Based Monitoring system provides a broad suite of tools to convey diverse real time and historical information from many sources on operation of the CMS detector. The services are made accessible both locally and remotely, to address the challenges of a large, geographically dispersed collaboration. The infrastructure also performs vital functions in the recording and quality determination of CMS event data. The WBM system has been a key element in the successful operation of the CMS experiment at CERN.

Acknowledgments

We congratulate our colleagues in the CERN accelerator departments for the excellent performance of the LHC and thank all CMS collaborators, the technical and administrative staffs at CERN and at other CMS institutes for their contributions to the success of the CMS effort. We acknowledge the enduring support for the development and operation of the CMS Web Based Monitoring effort provided by the funding agencies, CERN, the US Department of Energy, and the US National Science Foundation.

References

- [1] CMS Collaboration, *The CMS experiment at the CERN LHC*, 2008 JINST 3 S08004
- [2] Evans L (ed.) and Bryant P (ed.), *LHC Machine*, 2008 JINST 3 S08001
- [3] L. Tuura *et al.*, *CMS data quality monitoring: systems and experiences*, 2010 J. Phys.: Conf. Ser. 219 072020.
- [4] The CMS Collaboration, *The TriDAS Project, Technical Design Report, Volume 2: Data Acquisition and High Level Trigger*, CERN/LHCC 2002-26, 2002.
- [5] Arcidiancono R *et al.*, *CMS DCS Design Concepts*, 10th ICALEPCS Int. Conf. On Accelerator & Large Expt. Physics Control Systems, (2005) PO1.062-6.
- [6] Bauer G *et al.*, *Status of the CMS Detector Control System*, 2012 J. Phys.: Conf. Ser. 396 012023.
- [7] Lewis J, *et al.*, *The CERN Central Timing, a Vertical Slice*, Proc. International Conference on Accelerator and Large Experiment Control Systems, Knoxville, TN (2007) FOAA03.
- [8] Salter W *et al.*, "DIP Description", LHC Data Interchange Working Group (LDIWG) (2004), <https://edms.cern.ch/file/457113/2/DIPDescription.doc>
- [9] Gaspar C *et al.*, *DIM, a Portable, Light Weight Package for Information Publishing, Data Transfer and Inter-process Communication*, International Conference on Computing in High Energy and Nuclear Physics Padova, Italy, 1-11 February 2000.
- [10] Bauer G *et al.*, *Monitoring the CMS data acquisition system*, 2010 J. Phys.: Conf. Ser. 219 022042.
- [11] Cavallari F *et al.*, *Time-critical Database Condition Data Handling in the CMS Experiment During the First Data Taking Period*, 2011 J. Phys.: Conf. Ser. 331 042007
- [12] <http://www.oracle.com/us/products/database/options/active-data-guard/overview/index.html>
- [13] <http://scientificlinux.org>, <http://linux.cern.ch>
- [14] <http://www.apache.org>, <http://tomcat.apache.org>
- [15] Brun R and Rademaker R, ROOT: *An Object-Oriented Data Analysis Framework*, *Nucl. Inst. Meth. A* **389** (1997) 81-86. <http://root.cern.ch>
- [16] www.jfree.org/jfreechart
- [17] <http://jquery.com>
- [18] <http://www.highcharts.com>
- [19] <https://op-webtools.web.cern.ch/op-webtools/vistar/vistars.php?usr=LHCCMS>
- [20] <http://www.jcp.org/en/jsr/detail?id=286>
- [21] Jeitler M *et al.*, "The Upgrade of the CMS Trigger System", 2014 JINST 9 C08002
- [22] Hegeman J *et al.*, "The CMS Timing and Control Distribution System", 2015 Proc. IEEE Nucl. Sci. Symposium.