

TIPP 2011 - Technology and Instrumentation for Particle Physics 2011

CMS Online Web-Based Monitoring

William Badgett^a, Irakli Chakaberia^b, Juan Antonio Lopez-Perez^a, Kaori Maeshima^a, Sho Maruyama^{a,d}, Aron Soha^a, Balys Sulmanas^a, Zongru Wan^{a,b,c,*}

^aFermi National Accelerator Laboratory, Batavia, IL, 60510, USA

^bKansas State University, Manhattan, KS, 66506, USA

^cState University of New York at Buffalo, Buffalo, NY, 14260, USA

^dUniversity of California at Davis, Davis, CA, 95616, USA

Abstract

For large international High Energy Physics experiments, modern web technologies make the online monitoring of detector status, data acquisition status, trigger rates, luminosity, etc., accessible for the collaborators anywhere and anytime. This helps the collaborating experts monitor the status of the experiment, identify the problems and improve data taking efficiency. We present the online Web-Based Monitoring project of the CMS experiment at the LHC at CERN. The data sources are relational databases and various messaging systems. The project provides a vast amount of in-depth information including real-time data, historical trends and correlations in a user-friendly way.

© 2012 Published by Elsevier B.V. Selection and/or peer review under responsibility of the organizing committee for TIPP 11. Open access under [CC BY-NC-ND license](https://creativecommons.org/licenses/by-nc-nd/4.0/).

Keywords: Web-based monitoring; Online

1. Introduction

The Compact Muon Solenoid (CMS) [1] experiment is one of the high energy physics experiments at the Large Hadron Collider (LHC) [2] at CERN [3]. Currently it has about 3000 collaborators from 172 institutes spread over 39 countries. For such a large international collaboration, it is important to promptly relay the status of the experiment to the collaborators around the world. For this purpose, the CMS online Web-Based Monitoring (WBM) has been developed since early 2006.

The CMS online WBM server is accessible on web browsers for the collaborators locally and remotely, anywhere and anytime. Fig. 1 shows the main page. A suite of user-friendly tools is used widely to monitor the real-time and historical status of the experiment. The tools are heavily used by participants of the daily meetings of the CMS Commissioning and Run Coordination teams, to identify and understand problems, with the ultimate goal of improving data taking efficiency. WBM has become one of the most important collaborating tools for the experiment.

*Email: zrwan@fnal.gov

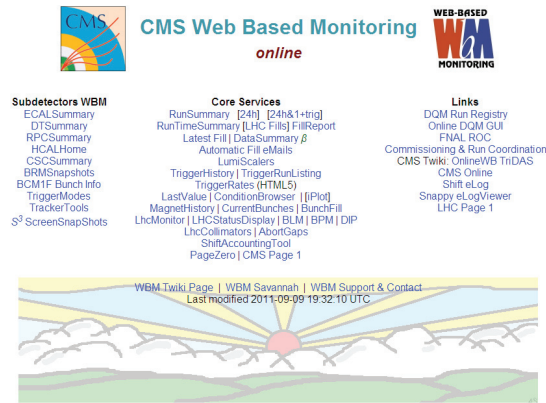


Fig. 1. CMS online Web-Based Monitoring main page.

2. Data sources

The data sources consist of the online Oracle database and various messaging system. The online Oracle database has a wealth of information for the real-time and historical states of the experiment including detector status, data acquisition status, trigger rates, luminosity, etc. The messaging systems include an LHC hardware serial line, an LHC software messaging system, CMS online real-time information, text messages from the shift leader, and others.

3. Software technologies

We aggregate multiple heterogeneous software and hardware data sources into one integrated user-friendly web interface. Modern software technologies are used to do so. While the format and transmission mechanisms for the source information is specific to the CMS experiment, large portions of the underlying infrastructure and software technology are compatible with being used elsewhere in the physics community.

On the server side, WBM runs on Scientific Linux [4] operating system and is based on the Apache/Tomcat application container [5] and the Java Servlet technology. JDBC and PL/SQL are used to read data from the database. A dedicated CERN java application called Data Interchange Protocol is used to receive data from LHC accelerator systems. A dedicated CMS application XDAQ flashlist [6] is used to receive data from CMS online systems. Plots for data visualization are mostly implemented by C++ and ROOT [7] with output in the formats with which high energy physicists are mostly familiar: *eps*, *png* and *root*. JFreeChart [8] is used to produce *png* plots directly in the Java Servlets in the cases where other formats are not required.

On the client side, data is presented in web pages using HTML, CSS and JavaScript supported universally by all of the web browsers. Users are not required to install any plugins. AJAX and XML are used to communicate with the server for the applications with an automatic refreshing feature. Recently, jQuery and HTML5 canvas [9] are also used for data visualization on the client side.

For the collaborative code development, SVN [10] is used for version control, TWiki [11] is used for documentation and Savannah [12] is used for task management. For code testing, we strive for one command to get things done: one command to set up the configuration, one command to compile and one command to install. Several development servers are used for testing. The services are unavailable for only a few seconds when a new release is pushed to production, thereby minimizing the interruption.

For CMS collaboration members, controlled access to the WBM server from the outside world is via CERN Single Sign On [13] authentication and a proxy to the private CMS domain.

4. WBM products

In the following, we describe some of the examples of the WBM products including the applications for real-time data, the applications for archival data and the applications for data taking efficiency.

4.1. Applications for real-time data

The applications for real-time data show the status of the experiment right now and they are automatically refreshed. A high-level non-expert type of page for general CMS members called CMS page 1 is shown in Fig. 2. It contains the information about the current LHC beam status, the current CMS data acquisition status, a dynamic plot of how much data is delivered by LHC and how much data is recorded by CMS during last 24 hours, the comments of the shift leader in the control room, etc. The page is made of a picture of 800×600 pixels in *png* format. It fits and works on almost all kinds of computers and devices.

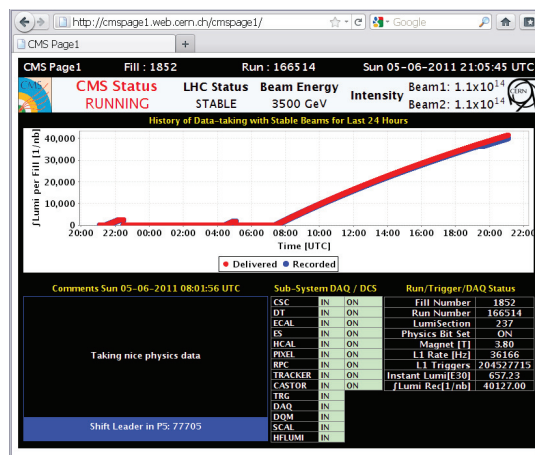


Fig. 2. CMS page 1.

We have applications to monitor the current status of LHC. One example for the LHC status is shown in Fig. 3. It displays the beam state in the context of possible states: beam preparations, stable beams, beam dump, etc. Another example for the bunch pattern of the LHC beams is shown in Fig. 4. A beam consists of bunches of particles with precisely configured gaps. The application displays the bunches of the beams as seen at CMS.

Most of the WBM real-time applications are for monitoring the current status of CMS. One example for the trigger rates is shown in Fig. 5. The application displays trigger configuration, rates and alarm. The time trend of each trigger path can be easily added or removed in the plot by toggling a check box. The plot is implemented by the new HTML5 canvas technology and works on the web browsers of both computers and phone devices. Another example for the status of luminosity sections is shown in Fig. 6. A luminosity section, which is 2^{18} orbits around the LHC tunnel with a speed very close to the speed of light and totals about 23 seconds, is the smallest quantized time range to define data quality. The application shows luminosity, beam status, high voltage status of sub-detectors, etc., for each luminosity section. This is a very important application for checking the online data quality. One more example is the slow control status shown in Fig. 7. It shows voltage, current, temperature, etc. for each sub-detector, displayed by image map corresponding to its geometry. A user can use a mouse to point to a component to display its values, click the component to show its statistics during last 24 hours, select sub-channels and plot time trends to diagnose problems.

The following is a concrete example of how the services were used to spot a problem, and how the acquired knowledge was fed back into actions. On August 20, 2010, the WBM trigger rate real-time monitoring tools showed that L1 jet rates periodically moved between two different values. Since the WBM

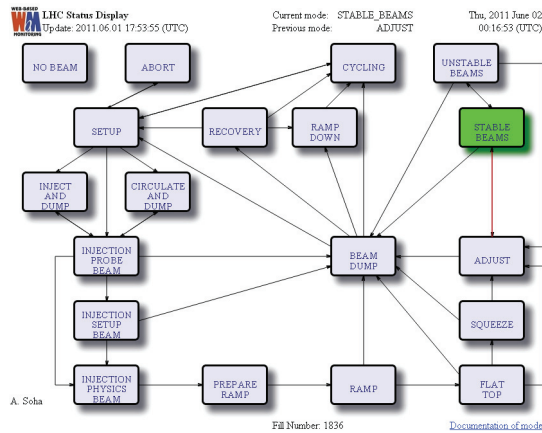


Fig. 3. LHC status display.




Fig. 4. Bunch pattern of the LHC beams.

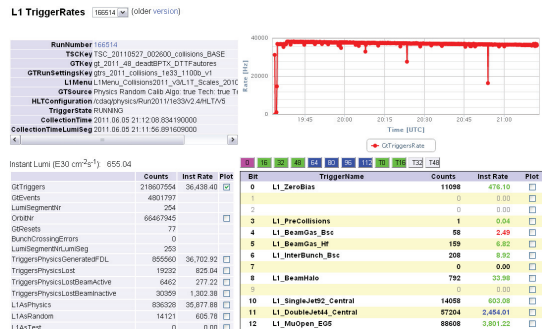


Fig. 5. Trigger rates.

tool showed that only jet triggers were affected, the shift crew figured out that the issue was related to the calorimeters. With this information in hand, the L1 trigger experts were called at 3am. It was found that some cables of the regional calorimeter system were malfunctioning, and the cables were replaced.

4.2. Applications for archival data

The applications for archival data in the online database are for the historical states. One example is called fill report, from which it is easy to drill down to the details. A fill is a period of time when beams




Fig. 6. Status of luminosity sections.

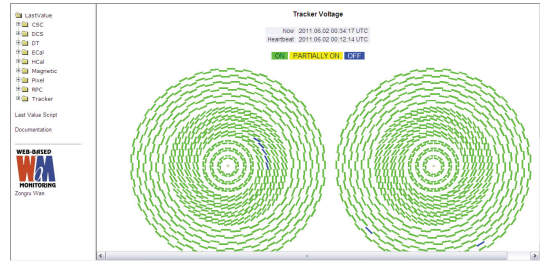


Fig. 7. Slow control status.

circulate in the LHC. At the front page, it shows a list of fills and summary information for each fill, shown in Fig. 8. If a user clicks one of the fills, it drills down to the CMS run numbers of the fill. A run is a period of time CMS takes data with a well defined configuration. If a user clicks one of the run numbers, it shows the detailed information of the run, including luminosity delivered and recorded, configurations, events, magnetic field, data acquisition components, level 1 and high level trigger rates, data file size, etc. If one wants to find out e.g. high level trigger rates, one can drill down further. Everything is amazingly linked together! In fact, this application is important for not only online monitoring but also offline data analysis.

Specific Fill: Begin: 1817 End: 1836 Stable: Last n Fills: 20 Stable: < 1835

Fill	CreateTime	Duration Stable	PeakInstLumi	Peak	PeakSpecLumi	DeliveredLumi	RecordedLumi	EmbyLumi	EmbyTime
	CMS Fill Declared	#Fills	$\mu\text{b}^{-1}\text{s}^{-1}$	μb^{-1}	μb^{-1}	μb^{-1}	μb^{-1}	%	%
1835	2011.06.01 16:48:10	1.25	1163.750	7.179	727.097	5643.718	3892.812	68.983	95.225
1835	2011.06.01 16:01:31	1.96	1148.200	7.090	728.459	4413.592	4489.536	99.959	100.000
1834	2011.06.01 15:04:31	0.000	0.000	0.000	0.000	0.000	0.000	0.000	100.000
1832	2011.06.01 11:46:46	0.000	0.000	0.000	0.000	0.000	0.000	0.000	100.000
1832	2011.06.01 07:31:04	0.000	0.000	0.000	0.000	0.000	0.000	0.000	100.000
1831	2011.06.01 04:59:28	0.000	0.000	0.000	0.000	0.000	0.000	0.000	100.000
1830	2011.06.01 00:22:33	0.000	0.000	0.000	0.000	0.000	0.000	0.000	100.000
1829	2011.05.31 20:29:18	0.000	0.000	0.000	0.000	0.000	0.000	0.000	100.000
1828	2011.05.31 17:48:07	0.000	0.000	0.000	0.000	0.000	0.000	0.000	100.000
1827	2011.05.31 15:13:26	0.000	0.000	0.000	0.000	0.000	0.000	0.000	100.000
1826	2011.05.31 14:19:40	0.000	0.000	0.000	0.000	0.000	0.000	0.000	100.000
1825	2011.05.31 13:46:02	0.000	0.000	0.000	0.000	0.000	0.000	0.000	100.000
1824	2011.05.31 08:43:34	0.000	0.000	0.000	0.000	0.000	0.000	0.000	100.000
1820	2011.05.31 04:22:31	1.12	1177.969	7.472	687.762	4690.250	3621.061	72.995	89.889
1822	2011.05.31 00:44:45	1.27	1266.741	7.819	709.845	6596.051	2655.986	42.766	46.723
1821	2011.05.30 21:19:33	0.000	0.000	0.000	0.000	0.000	0.000	0.000	100.000
1820	2011.05.30 19:29:49	0.000	0.000	0.000	0.000	0.000	0.000	0.000	100.000
1819	2011.05.30 17:11:21	0.000	0.000	0.000	0.000	0.000	0.000	0.000	100.000
1818	2011.05.30 13:58:19	0.000	0.000	0.000	0.000	0.000	0.000	0.000	100.000
1817	2011.05.30 09:19:25	0.000	0.000	0.000	0.000	0.000	0.000	0.000	100.000
total		5.10	1266.741	7.819	727.097	21081.872	14429.896	68.765	82.453

Fig. 8. Fill report

The instantaneous luminosity scenario for trigger design has changed from 10^{31} to $5 \times 10^{33} \text{ cm}^{-2}\text{s}^{-1}$, a factor of 500, since last year. It is an enormous challenge for the trigger designers to deal with the rapidly increasing instantaneous luminosity. We provide an application to monitor the trend of trigger cross sections versus instantaneous luminosity, shown in Fig. 9 for one of the level 1 trigger paths. The parameter fitting and the visualization for the linear and non-linear effects for all of the trigger paths have been a great help for designing the trigger tables.

A large amount of data in the online database is in the form of values and time stamps. We provide a generic application called condition browser to visualize such data. Time trends can be plotted in any given




Fig. 9. Trigger cross section versus instantaneous luminosity.

time range. The implementation starts with meta data (where to find the values and the time stamps) and automates the rest of the process. This design makes it very flexible to implement requests e.g. “I want this and that during the last hour by one single click,” instead of writing any specific combinations from scratch. There are thousands of variables in the condition browser. An example for two variables versus time put together is shown in Fig. 10 (a). Furthermore, correlations between any two variables can be joined by closest time difference. This is fondly called at CMS “anything vs. anything.” An example correlation for the two variables in Fig. 10 (a) is shown in Fig. 10 (b). Again, work is reduced to tell the condition browser where to find the two variables and the rest is automated. This design makes it extremely easy to correlate two arbitrary asynchronous variables no matter where they are stored. The output is implemented in various formats including *png*, *eps*, *html*, *txt*, *csv*, *xml*, *sql*, *root* file and *root* script.

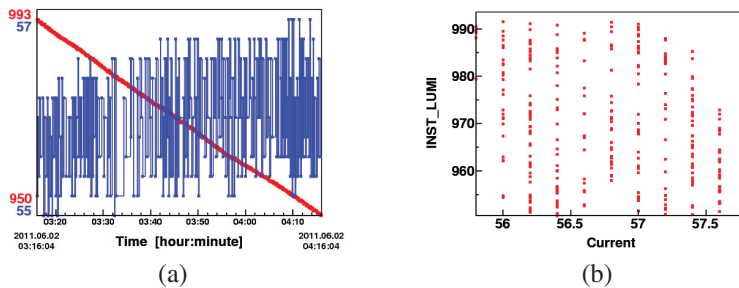


Fig. 10. Condition browser. (a) value versus time for two variables put together; (b) correlation of two variables joined by closest time difference.

4.3. Applications for data taking efficiency

The CMS experiment needs very high data taking efficiency in order to stay competitive. To help identify problems and improve data taking efficiency, we provide an application called run time logger, shown in Fig. 11. The shift leader in the control room is required to categorize, and identify the reasons for, the time lost and luminosity lost. To quantify data taking efficiency, we provide an application for data taking efficiency versus fill, shown in Fig. 12. The red bars are for efficiency by time and the blue bars are for efficiency by luminosity. These applications are very important for the operation of the CMS experiment.

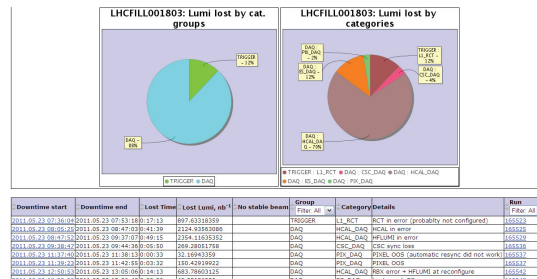


Fig. 11. Run time logger.




Fig. 12. Data taking efficiency versus fill.

5. Summary

The CMS online Web-Based Monitoring is accessible to the collaborators locally and remotely, anywhere and anytime. It provides a vast amount of in-depth information including real-time data, historical trends and correlations in a user-friendly way. It is a key element for successful data taking operation of the CMS experiment.

Acknowledgements

We would like to thank all of our CMS colleagues for their contributions to this work and the funding agencies for their support.

References

- [1] Compact Muon Solenoid: <http://cms.cern.ch>
- [2] Large Hadron Collider: <http://lhc.web.cern.ch/lhc>
- [3] CERN: <http://www.cern.ch>
- [4] Scientific linux at CERN: <http://linux.web.cern.ch/linux>
- [5] Apache/Tomcat: <http://httpd.apache.org/>, <http://tomcat.apache.org/>
- [6] G. Bauer et al., “Monitoring the CMS data acquisition system”, 2010 *J. Phys.: Conf. Ser.* 219 022042 <http://dx.doi.org/10.1088/1742-6596/219/2/022042>
- [7] ROOT: <http://root.cern.ch/>
- [8] JFreeChart: <http://www.jfree.org/jfreechart>
- [9] HTML5 canvas: <http://www.whatwg.org/specs/web-apps/current-work/multipage/the-canvas-element.html>
- [10] SVN: <http://subversion.apache.org/>
- [11] TWiki: <http://www.twiki.org/>
- [12] Savannah at CERN: <https://savannah.cern.ch/>, based on Savane: <http://gna.org/projects/savane>
- [13] CERN Single Sign On: <https://cern.ch/authentication>