Towards a centralized Grid Speedometer

I Dzhunov[1], J Andreeva[1], E Fajardo[2], O Gutsche[3], S Luyckx[4], P Saiz[1]

[1]CERN, European Organization for Nuclear Research, Switzerland
[2]Universidad de los Andes, Colombia
[3]FNAL, Fermi National Accelerator Laboratory, USA
[4]University of Antwerp, Belgium

idzhunov@cern.ch

Abstract. Given the distributed nature of the Worldwide LHC Computing Grid and the way CPU resources are pledged and shared around the globe, Virtual Organizations (VOs) face the challenge of monitoring the use of these resources. For CMS and the operation of centralized workflows, the monitoring of how many production jobs are running and pending in the Glidein WMS production pools is very important. The Dashboard Site Status Board (SSB) provides a very flexible framework to collect, aggregate and visualize data. The CMS production monitoring team uses the SSB to define the metrics that have to be monitored and the alarms that have to be raised. During the integration of CMS production monitoring into the SSB, several enhancements to the core functionality of the SSB were required; They were implemented in a generic way, so that other VOs using the SSB can exploit them. Alongside these enhancements, there were a number of changes to the core of the SSB framework. This paper presents the details of the implementation and the advantages for current and future usage of the new features in SSB.

1. Towards a centralized Grid Speedometer

Virtual Organizations (VOs) face the challenge of effectively using the CPU resources of the Worldwide Large Hadron Collider (LHC) [1] Computing Grid (WLCG) [2]. To tackle this challenge, monitoring the use of these resources becomes a central part of the computing operations. The CMS experiment [3], as one of the four main VOs operating at the LHC, is facing this challenge as well. Moreover the group of experts who form the CMS computing operations team are also distributed worldwide [4]. To achieve this task, the team makes use of information on the number of running (Figure 1) and pending jobs at each site, correlated with the status of the sites in order to create alarms for underperforming sites. In addition, important information such as the number of production slots at sites is stored in text files. Maintaining and accessing all those sources of data is difficult. Thus a single, authoritative, widely accessible monitoring tool was needed to gather and display all information. The Site Status Board (SSB) matched all these requirements and was adopted by CMS for monitoring their usage of the different sites by their production infrastructure.
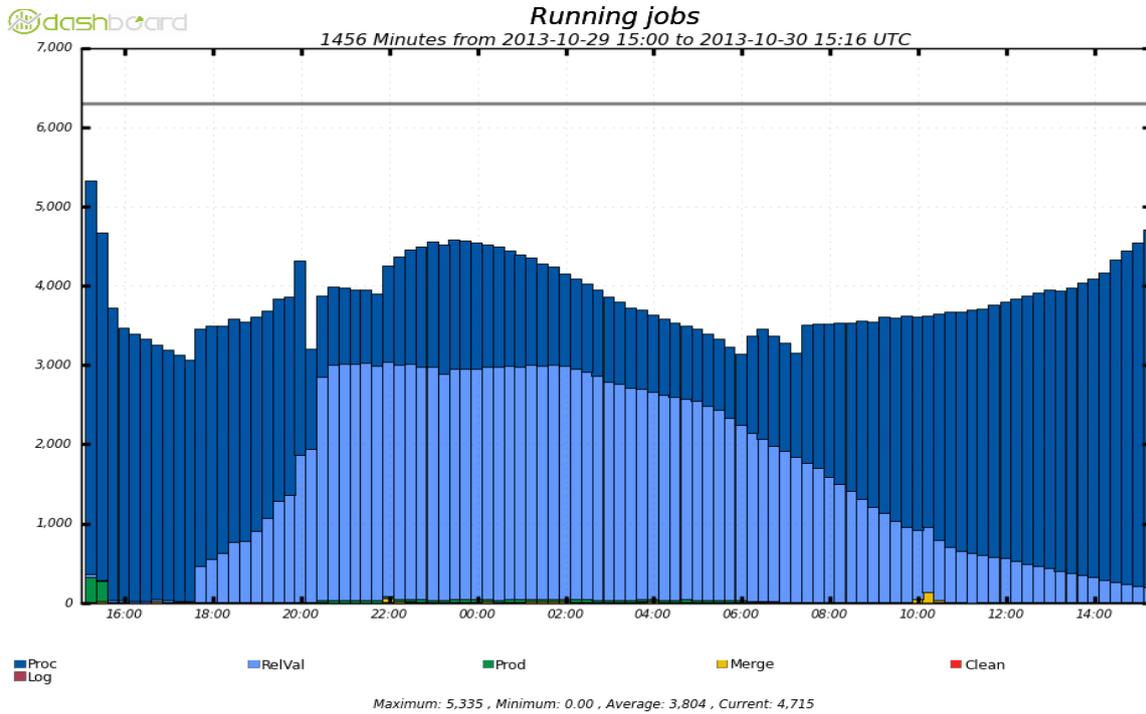
Figure 1. SSB plot showing number of production running jobs at a site including pledged slots.

2. Introduction to Site Status Board

The SSB, built within Experiment Dashboard [5] framework, is a highly customizable and easily configurable monitoring tool that is a container for metrics and can provide their visualization (both current overview and evaluation of historical data). Metrics are defined by the VOs as the measurement of site performance or result of a functional test. Logically related metrics can be aggregated into views so that VOs can refer to a quick overview of the health of the distributed infrastructure. VOs can define as many views as necessary.
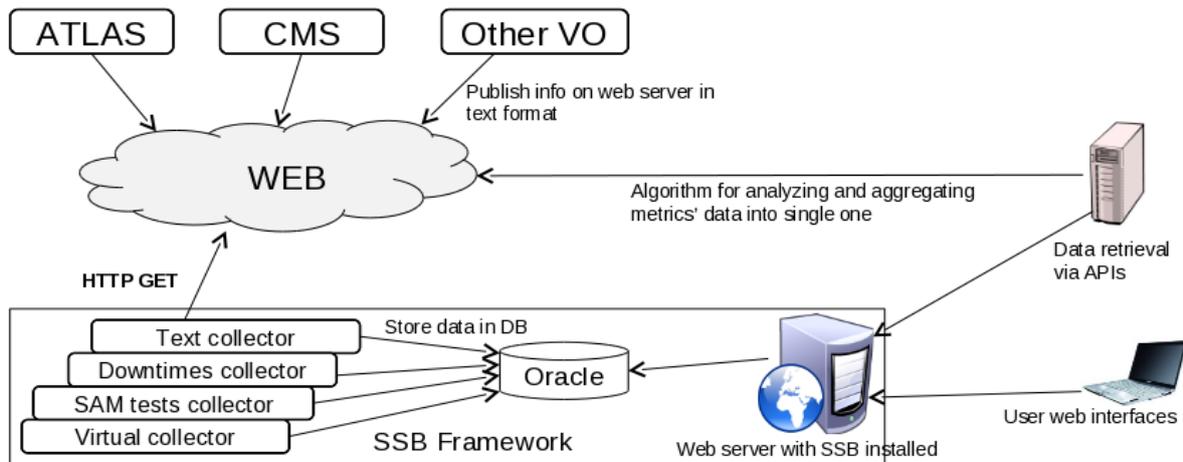

Figure 2. Data flow of SSB.

VO communities can decide what data they want to monitor and publish it on the web in text format. The SSB offers a text collector which gets the data from a known URL and stores it in an Oracle database. As depicted in Figure 2, stored data can be retrieved in various formats from the SSB web server. The end user can view data in HTML pages with high quality historical plots as well as retrieve data via APIs. SSB provides data access in HTML, JSON and XML format.

In addition, the SSB offers some common collectors, such as the Downtimes collector which collects downtime information for sites, SAM tests [6] collector that collects results from SAM tests, and Virtual collector offering combination of basic metrics via the logical AND operator. Metrics collected in SSB can be further processed by external custom applications and the result of this processing fed back to SSB as new metrics. This functionality is widely used by the LHC community and one example is the CMS Site Readiness [7] metric. The objective of this metric is to summarize for each day if a given site is ready to receive both analysis and production workflows. It is calculated using input from the SAM tests, Hammer Cloud tests [8], PhEDEx [9] network transfer links availability and site downtime information. Then the newly calculated metric "Site Readiness" is uploaded back to the SSB.

3. Recent changes in SSB
This section will present some recent enhancements to SSB. The CMS production team already exploits some of these and there are others from which it could be profited in the future.

3.1. Metric data modification
Sometimes data providers publish inaccurate data (i.e. a downtime was not taken into account when calculating the Site Readiness metric). Thus the ability to modify metric data was required, so that plots and reports based on SSB data are accurate. As depicted in figure 3, the SSB team has introduced the ability to modify metric data (historical and current) from the user interface. Only users with special privileges can do it and only for metrics that allow data modifications. This feature was adopted in the CMS production view for keeping track of the number of production pledges and overall slots at sites. Historically, such data were stored in text files but a solution which saves, modifies and displays them in a central place, such as SSB is clearly preferable.
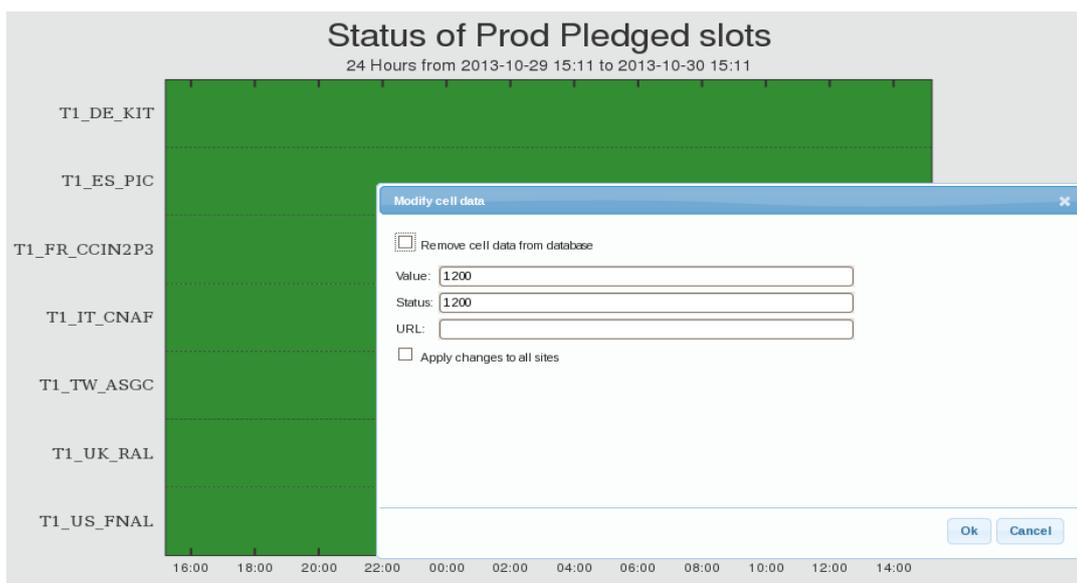

Figure 3. Modify SSB metric data.

3.2. Client side caching of data for plots

In the initial implementation the SSB queried data from the database every time a user requested a plot. The new implementation allowed to improve the user experience and lighten the database load. Currently the SSB caches data from the last database call on the client and, if the data for the next search can be retrieved from the recently cached data, a database call is avoided. Data can be taken from the cache if new search filters are more specific than the ones applied when data was cached (i.e. filtering from all sites to particular site or reducing the time window from one week to 48 hours). If the new search filters are more general, then data is retrieved from the database and the client cache is updated with the more general data. This improvement was implemented for the historical plots and all metrics (including ones from the CMS production view) benefit from it.

3.3. Compression of collected data

As the number of metrics monitored by SSB grows, the stored data volume increases and queries for retrieving data become correspondingly slower. SSB developers are facing the challenge of how to store data effectively and how to retrieve it optimally. Instead of keeping raw data, SSB stores only status changes. This is appropriate for status metrics, where transition from one status to another is infrequent, but inappropriate in the cases of numeric metrics (i.e. number of running jobs) where numbers are changing every few minutes. For those cases compression algorithms have been implemented in SSB. Privileged users can configure which metric data is to be compressed, the data (i.e. older than x days) to be compressed and the compression algorithm to be used (i.e. daily average). In this manner, data for some metrics can be compressed up to a factor of 240 (replacing 10 min intervals with one daily averaged interval) and retrieval and plotting of historical data is much faster. The number of running and pending jobs in the production view are perfect candidates to benefit from these compression algorithms.
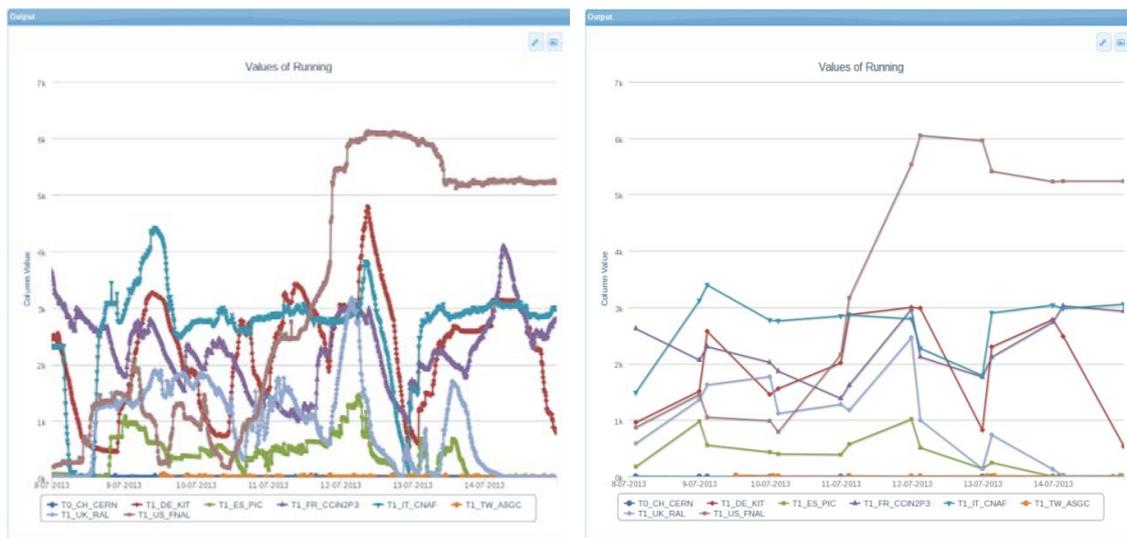
Figure 4. Retrieval and plotting of compressed data is faster.

3.4. Alarms

As a modern monitoring tool, SSB has seen the need of notifying users and developers for unusual circumstances.

3.4.1. Alarms to site responsible people on status change. Authorized SSB users can register to receive alarms when a site changes color/status in a particular view. The color in a view is defined as 'worst color from the colors of the critical metrics in a view'. As a view may consist of a single metric,

this mechanism can also be used for notifications when a site changes color/status in a single metric. There is no need for the site system administrators to refresh a web interface periodically to see if a site is working properly; SSB does this automatically and sends an email with a detailed description of which site has changed color (including the specific change details) in which view, which metric caused the change of color, when the change occurred and a link to the SSB site history plot for the specified site and view. Over 25 site administrators are using this feature in CMS and production view users can benefit from it for the alarm metrics.

3.4.2. Alarms to metric providers when there is a problem with a data source. The integrity of collected data is very important in the WLCG world. In SSB it is observed that user communities have problems with data publishers. If those problems are not spotted on time, data would be lost for long periods. For this reason, SSB developers recently implemented a mechanism to notify developers and metric providers if the source of a metric becomes unavailable or stops updating. Alarms are sent as soon as a problem with a metric source arises and this alarm contains a reminder list of those metrics whose sources are currently not being updated. Another alarm is sent when the source of a metric resumes being updated. Those alarms are largely used by ATLAS and can be enabled for CMS as well.

4. Summary
Straight forward integration of CMS production monitoring into the SSB has demonstrated that SSB is a generic monitoring tool which can be effectively used for various computing tasks performed by the LHC computing community. As a result of this integration, SSB became a key factor in the monitoring of the CMS production workflows. The CMS distributed production team and site administrators use the tool to determine whether sites are running production tasks at full speed and to debug eventual problems with the production servers. Apart of monitoring of the production workflows, CMS uses SSB for distributed computing shifts and site commissioning. Currently CMS is considering the use of SSB for monitoring distributed analysis and analysis servers. The rich functionality and flexibility that SSB offers continues to attract more and more users and strengthens the role of SSB as one of the core systems for WLCG infrastructure monitoring.

References
[1]   Evans L and Bryant P, "LHC Machine", 2008, JINST 3 S08001 doi:10.1088/1748-0221/3/08/S08001
[2]   Bird I, "Computing for the Large Hadron Collider", 2011, Annual Review of Nuclear and Particle Science, 61:99–118 doi:10.1146/annurev-nucl-102010-130059
[3]   CMS Collaboration R. Adolphi et al., The CMS experiment at the CERN LHC, JINST, 0803, S08004 (2008)
[4]   Fajardo E et al "A new era for central processing and production in CMS ''" 2012J., Phys.: Conf. Ser. 396 doi:10.1088/1742-6596/396/4/042018
[5]   J. Andreeva et al, Experiment Dashboard for monitoring of the LHC distributed computing systems, 2011, J. Phys.: Conf. Ser. 331 072001 doi:10.1088/1742-6596/331/7/072001
[6]   Collados, D., Shade, J., Traylen, S., Imamagic, E. Evolution of SAM in an enhanced model for Monitoring WLCG services, Journal of Physics:Conference Series, Volume 219, Issue 6, pp. 062008 (2010).
[7]   J Flix et al 2011 J. Phys.: Conf. Ser. 331 072020 doi:10.1088/1742-6596/331/7/072020
[8]   Daniel C van der Ster et al 2011 J. Phys.: Conf. Ser. 331 072036 doi:10.1088/1742-6596/331/7/072036
[9]   Magini N, et al "Distributed data transfers in CMS"2011 Journal of Physics: Conference Series 331 042036 URL http://stacks.iop.org/1742-6596/331/i=4/a=042036