

Site Status Board - a flexible monitoring system developed in close collaboration with user communities

Julia Andreeva

*CERN, European Organization for Nuclear Research
CH-1211, Genève 23, Switzerland
E-mail: Julia.Andreeva@cern.ch*

Simone Campana

*CERN, European Organization for Nuclear Research
CH-1211, Genève 23, Switzerland
E-mail: Simone.Campana@cern.ch*

Alessandro Di Girolamo

*CERN, European Organization for Nuclear Research
CH-1211, Genève 23, Switzerland
E-mail: ale@diggi.org*

Ivan Dzhunov¹

*CERN, European Organization for Nuclear Research
CH-1211, Genève 23, Switzerland
E-mail: Ivan.Antoniev.Dzhunov@cern.ch*

Josep Flix

*PIC
Barcelona, Spain
E-mail: jflix@cern.ch*

Stavro Gayazov

*Budker Institute of Nuclear Physics
Novosibirsk, Russia
E-mail: stavro.gayazov@cern.ch*

Edward Karavakis

*CERN, European Organization for Nuclear Research
CH-1211, Genève 23, Switzerland
E-mail: Edward.Karavakis@cern.ch*

¹ Speaker

Lukasz Kokoszkiewicz

*CERN, European Organization for Nuclear Research
CH-1211, Genève 23, Switzerland
E-mail: Lukasz.Kokoszkiewicz@cern.ch*

Peter Kreuzer

*Rheinisch-Westfaelische Tech. Hoch.
Aachen, Germany
E-mail: pkreuzer@cern.ch*

Michal Nowotka

*CERN, European Organization for Nuclear Research
CH-1211, Genève 23, Switzerland
E-mail: Michal.Maciej.Nowotka@cern.ch*

Pablo Saiz

*CERN, European Organization for Nuclear Research
CH-1211, Genève 23, Switzerland
E-mail: Pablo.Saiz@cern.ch*

Jaroslava Schovancova

*CERN, European Organization for Nuclear Research
CH-1211, Genève 23, Switzerland
E-mail: Jaroslava.Schovancova@cern.ch*

Andrea Sciaba

*CERN, European Organization for Nuclear Research
CH-1211, Genève 23, Switzerland
E-mail: Andrea.Sciaba@cern.ch*

David Tuckett

*CERN, European Organization for Nuclear Research
CH-1211, Genève 23, Switzerland
E-mail: David.Tuckett@cern.ch*

Collaborative development proved to be a key of the success of the Site Status Board (SSB) which is heavily used by the experiments at the Large Hadron Collider (LHC) at CERN for the computing shifts and site commissioning activities. SSB is a tool that helps Virtual Organizations (VOs) to monitor the performance of sites. The selection, criticality and combination of monitoring metrics fall clearly in the domain of the VO administrators. Therefore VO administrators and computing teams define monitoring metrics and custom views of the monitoring data, in addition to developing sensors and data publishers. The responsibilities of the SSB team include development and support of the SSB framework and the SSB services which store, aggregate and visualize monitoring data. The collaboration extends beyond the customization of metrics and views to the development of new functionality and visualizations. SSB Developers and VO administrators cooperate closely to ensure that requirements are met and, wherever possible, new functionality is pushed upstream to benefit all users and VOs.

The contribution covers the evolution of SSB over recent years to satisfy diverse use cases through this collaborative development process.

*EGI Community Forum 2012 / EMI Second Technical Conference,
Munich, Germany
26-30 March, 2012*

1. Motivation for creation of such a tool

Site Status Board (SSB) developers want to help the experiments at the Large Hadron Collider (LHC) at CERN to find the answer to the question “How to define if a site is working properly for a Virtual Organization (VO)?”. Moreover sites also need to know how they are serving different experiments.

For this purpose a metric is defined as a measurement of site performance or result of a functional test. A metric’s value can be a number or a string (status). Metrics are defined by the VOs depending on what they want to measure. Logically related metrics are aggregated into views so that VOs can get a quick overview of the health of the distributed infrastructure. VOs can combine different metrics into as many views as needed.

2. The SSB Solution

The SSB solution, built within Dashboard Experiment [1] framework, includes:

- Flexible system that allows monitoring of various metrics with different frequency of measurements, validity of data, etc
- Visualization of metrics (current overview and historical data)
- Easily customized container for metrics

VOs participate in defining the functionality and design of the user interface. Their contribution is mandatory for the success of the system.

3. Data flow of the system

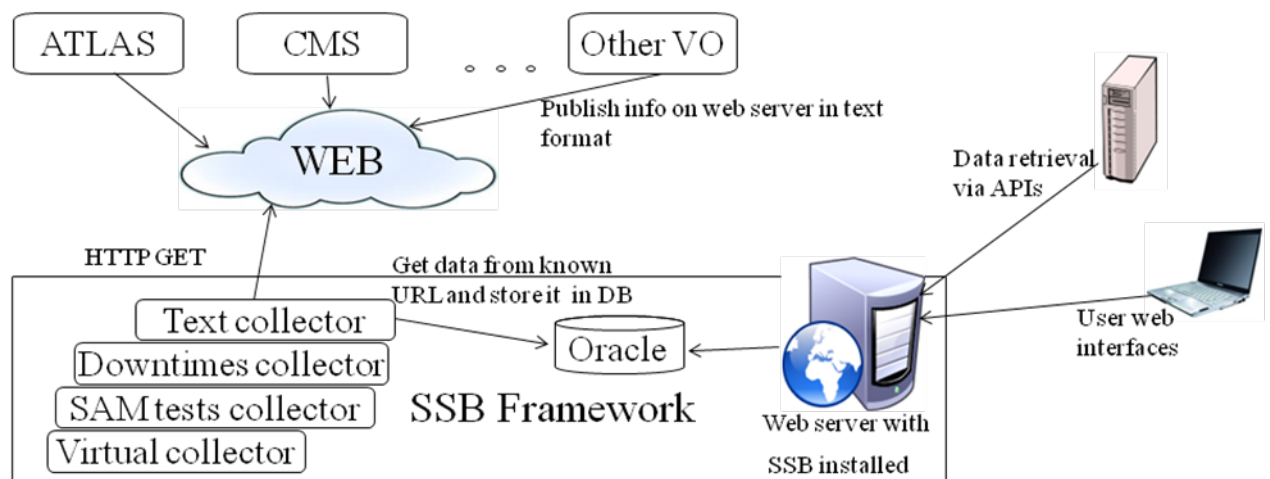


Figure 1 Data flow of SSB

Different VOs decide what data they want to monitor and publish it on the web in text format. The text collector which is part of the SSB framework gets the data from a known URL (where data is published for a specific metric) and stores it in an Oracle database. As depicted in figure 1, 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 in HTML, JSON and XML format.

A common API was agreed between SSB core developers and VO experts. The decoupling of work and responsibilities in the SSB workflow is one of the key elements of the success of SSB. VO experts provide the data, while SSB core developers retrieve, store and provide it back to the user in different formats. The SSB framework has some common collectors such as Downtimes collector that collects downtime information for sites, SAM tests [2] collector that collects results from SAM tests, and Virtual collector that offers basic metrics combination via the logical AND operator. VO experts can also retrieve metrics' data via API and use any other algorithm to combine metrics' data into a single metric which is fed back into SSB. CMS use this approach for Site Readiness calculations.

4. SSB development and responsible groups

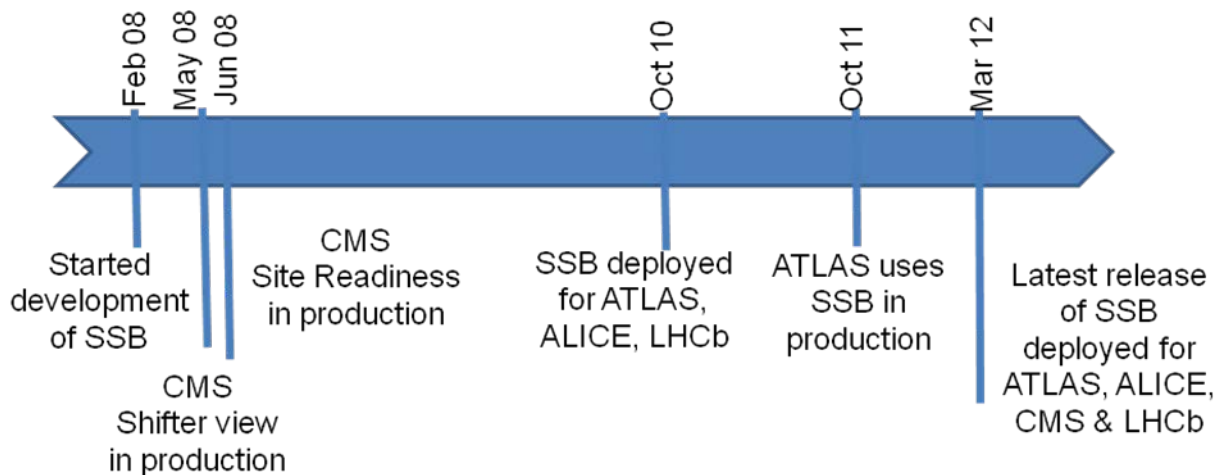


Figure 2 SSB Time line

The development of SSB started in February 2008. Only 4 months later, CMS were using their Shifter view and Site Readiness in production. In October 2010, thanks to the generic idea of SSB, the system was deployed for ATLAS, ALICE and LHCb. One year after that, ATLAS started using SSB in production. It took them one year to understand which metrics were important and should be monitored. So far ALICE and LHCb have used SSB only for testing purposes creating minimum set of metrics and views. Currently both experiments are still using the native monitoring systems. The latest release of SSB was deployed in March 2012. All these major SSB events are depicted in figure 2.

There are 3 categories of users and developers working with SSB:

- SSB core developers who develop common functionality
- VO experts who define metrics (what is the source URL of the data, how often measurements should be taken, what is the validity interval of the measurements)

and setup views. VO experts need write-access to the application, because they configure how the SSB “container” will be filled.

- End users - mostly site administrators and computing operators. End users have read-access to the application and they check the views defined by the VO experts.

5. How to set up new SSB

Setting up a new SSB instance is straightforward. First of all, the framework has to be installed. This includes installing a set of RPMs, configuring an Oracle database (for the time being the SSB DAO layer works only with Oracle), and setting up the SSB collectors. Once the framework is in place, the second step is defining the metrics. For each metric, monitoring information needs to be published on the web in text format. The format of the text files is - <datetime> <sitename> <status> <color> <link> . After a metrics’ data are provided on the web, metrics should be configured in SSB. Important configuration options are: URL from where data will be retrieved, frequency of data retrieval from this URL, validity interval of the data, etc. The final step is to aggregate logically related metrics into views.

It is up to the VO to define the metrics and views that need to be monitored. For instance, ATLAS uses the SSB to monitor quite a lot of different activities: data transfers, data subscriptions, site availability, status of particular services, some network measurements, etc. In total, they have more than thirty views, and around two hundred different metrics.

ATLAS feedback is that setting up new metrics and combining them into a view is very easy. They managed to do this in one day.

6. What SSB is used for

There are two main use cases of SSB: investigation of a problem and retrieving historical data

6.1 Investigation of a problem with SSB

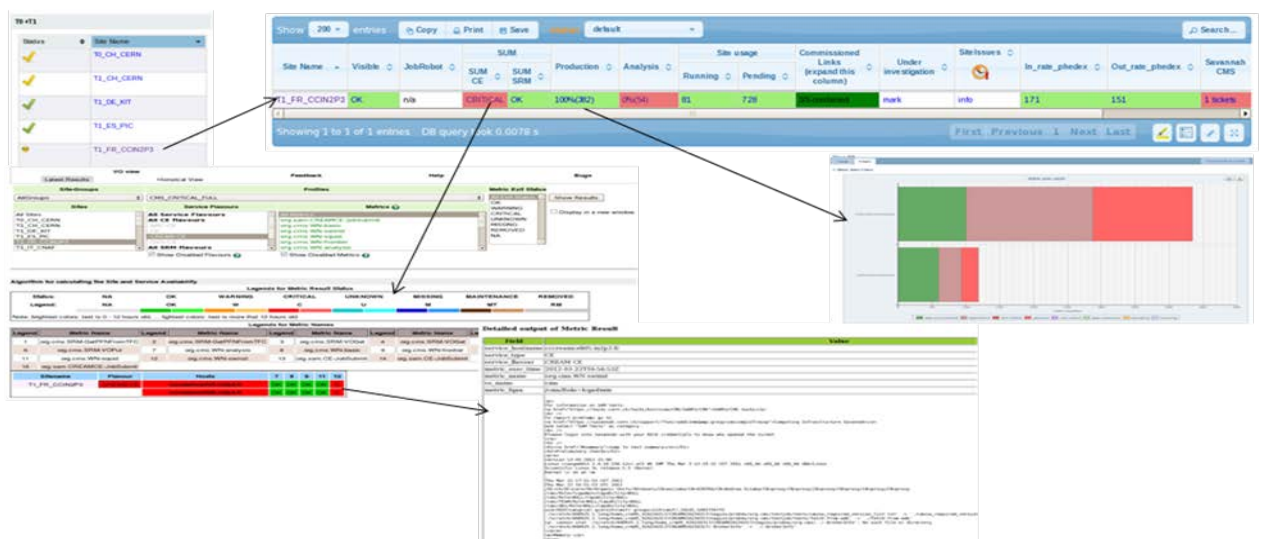


Figure 3 SSB – investigation of a problem

SSB shows the current situation at sites and has links to different applications for further information and digging into the problem (Figure 3). The index page of SSB displays the current status of sites. A green tick next to a site name means the site is working fine, a yellow tick means at least one of the critical metrics is in warning, a red dot means at least one of the critical metrics is failing, an “at work” symbol means at least one of the critical metrics reports maintenance color (brown), and a “lens” symbol means the site is under investigation (i.e. a ticket has been submitted in the Savannah ticketing system to follow up the problem). Clicking on a site name leads to another page in SSB where the current value/status of the metrics can be seen. Every metric’s cell is colored, shows the last value of the metric’s measurement, and has a link for further information. These links lead to different monitoring applications.

6.2 Historical data retrieval per metric to check how sites are performing over time

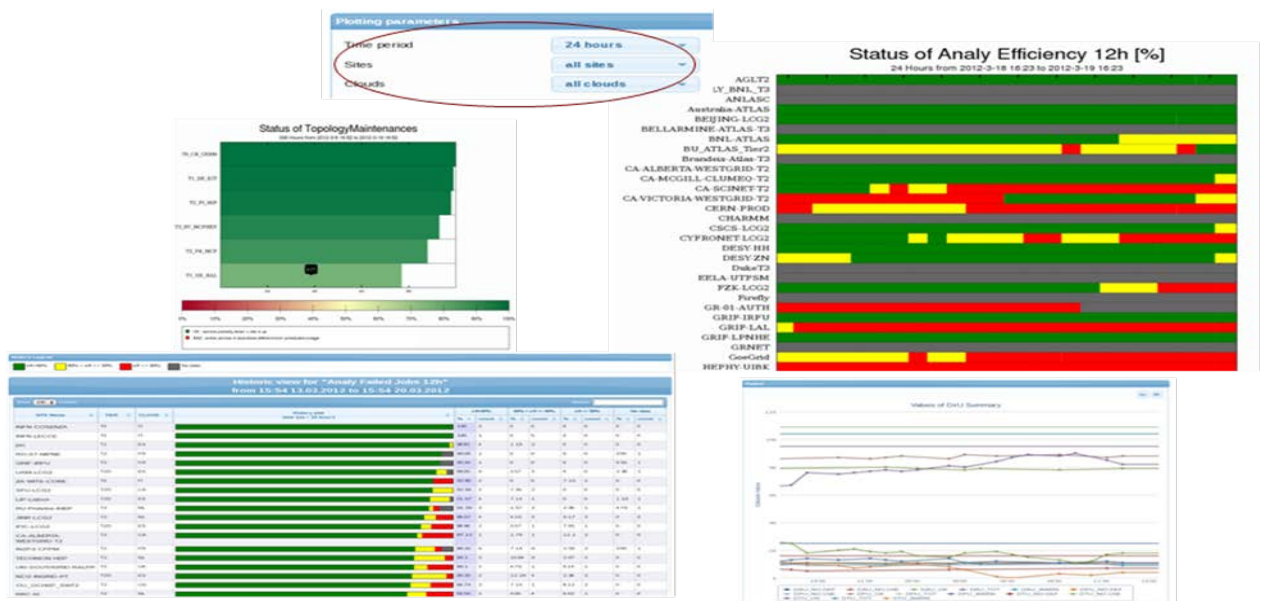


Figure 4 SSB – historical data retrieval

SSB offers plenty of historical plots that show how sites perform over time (Figure 4). Some plots show how sites changed from one status/color/value to another over time. Another plot shows what percentage of time sites spent in each status/color. Ranking plots were introduced recently into SSB. One ranking plot shows the percentage of time sites spent in “good” over “good” plus “bad” statuses. Another ranking plot shows the number of sites that were in each status.

Similar historical plots are available per site (showing a selection of metrics).

Historical data retrieval can be done not only in HTML format (figure 4), but also in XML and JSON.

7. VOs contributing to SSB core

The experiments that use SSB have already contributed to the development of the core. Thanks to this, all SSB users can benefit from the improvements made by other users. The next subsections describe two examples of these contributions.

7.1 ATLAS's contribution

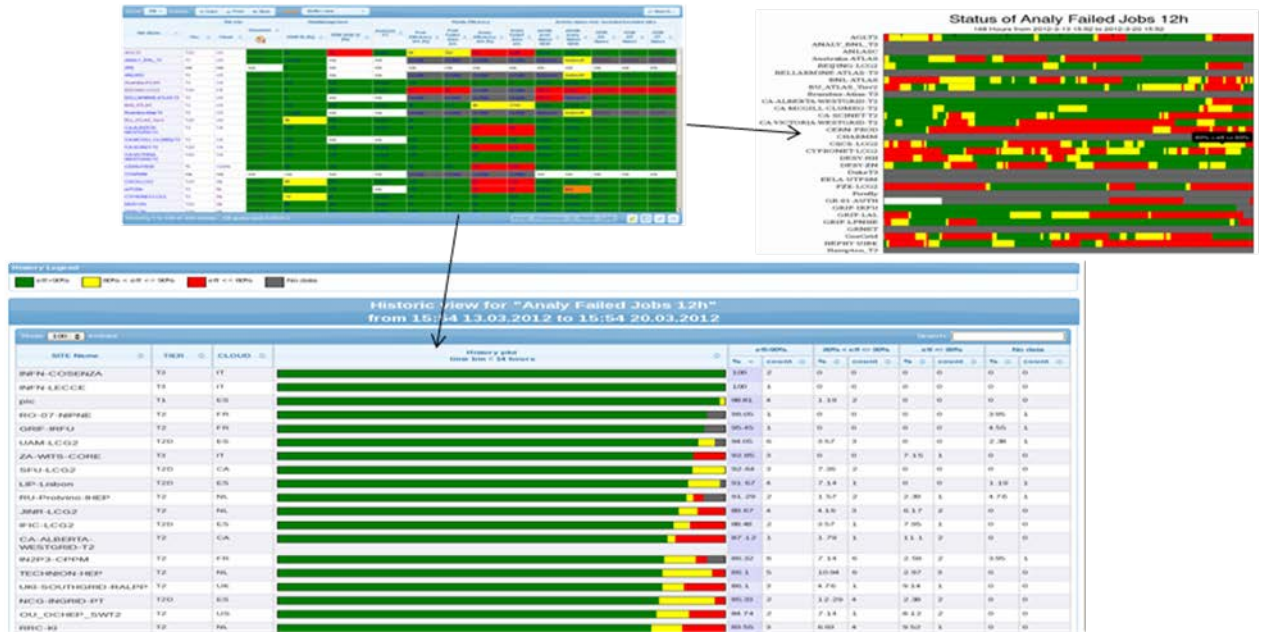


Figure 5 The bottom plot was implemented by ATLAS

Going from the Shifter view to metric history view, ATLAS people understood that they need more functionality. They implemented another metric history view with statistics (bottom plot on figure 5), from which they can rank site performance. Thanks to the core principle of SSB, this view is available for all other user communities.

7.2 CMS’s contribution



Figure 6 Bottom right plot is the new ranking plot in SSB

CMS people used to use another application to assess if a site is ready to be used (top right plot on figure 6). Data for this application is retrieved from Site Readiness Status metric in SSB. SSB core developers implemented this kind of ranking plot (bottom right plot on figure 6) and made it available for all metrics for all user communities. CMS are switching now to using the SSB plot for site readiness matters.

8. Siteview - another application using SSB core

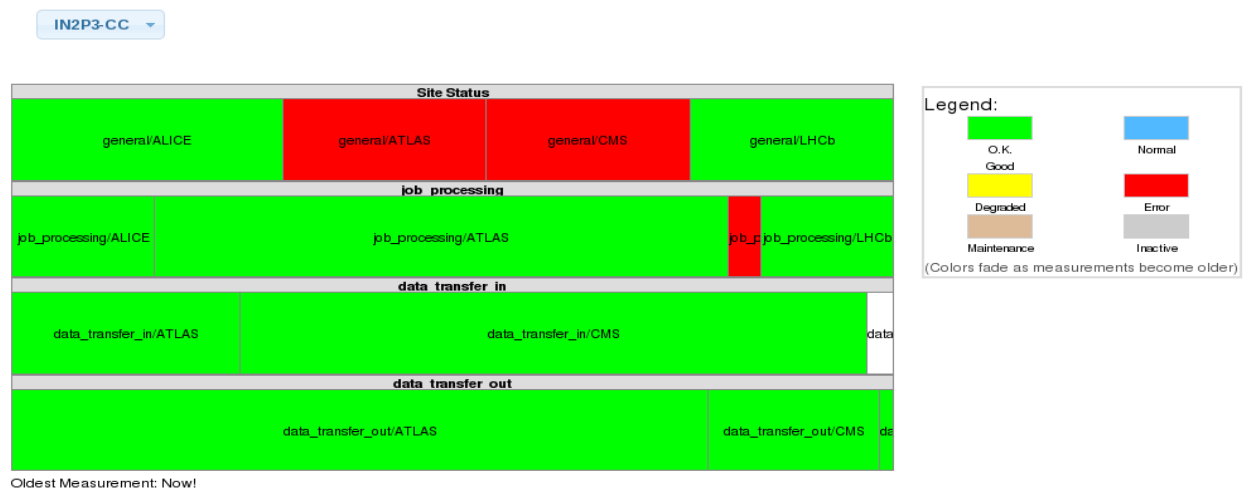


Figure 7 Siteview application

Siteview is another application that uses the SSB core. This application shows how sites are serving different experiments. Data for this application is provided by the VOs (again published on the web in text format). This application can be accessed here <http://dashb-siteview.cern.ch>

9. Conclusions

Close collaboration between SSB developers and the computing community of the LHC VOs has led to the creation of a successful, highly customizable, easy to configure and flexible solution for monitoring computing activities at distributed sites.

The SSB framework was designed in a generic way which allows it to be easily adapted to the needs of different Virtual Organizations.

CMS SSB is visited weekly by 220 unique visitors from 27 countries, CMS experts have defined 100 metrics and combined them into 10 views. ATLAS SSB is visited weekly by 133 unique visitors from 18 countries. ATLAS experts have defined 200 metrics and combined them into 30 views. You can access CMS and ATLAS SSB applications here : dashb-ssb.cern.ch and dashb-atlas-ssb.cern.ch .

References

- [1] 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](https://doi.org/10.1088/1742-6596/331/7/072001)
- [2] 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).