

PAPER • OPEN ACCESS

A Dashboard for the Italian Computing in ALICE

To cite this article: D Elia *et al* 2017 *J. Phys.: Conf. Ser.* **898** 092054

View the [article online](#) for updates and enhancements.

Related content

- [Time-series pattern recognition with an immune algorithm](#)
I Paprocka, W M Kempa, C Grabowik et al.
- [Critical services in the LHC computing](#)
A Sciabà
- [Measuring and understanding computer resource utilization in CMS](#)
J Andreeva, S Belforte, S Blyweert et al.

A Dashboard for the Italian Computing in ALICE

**D Elia¹, G Vino^{1,2}, S Bagnasco³, A Crescente⁴, G Donvito¹, A Franco¹, S Lusso³,
D Mura⁵, S Piano⁶, G Platania⁷**

on behalf of the ALICE Collaboration

¹ INFN Bari, Italy

² Consortium GARR Roma, Italy

³ INFN Torino, Italy

⁴ INFN Padova, Italy

⁵ INFN Cagliari, Italy

⁶ INFN Trieste, Italy

⁷ INFN Catania, Italy

E-mail: Domenico.Elia@ba.infn.it, Giacchino.Vino@ba.infn.it

Abstract. A dashboard devoted to the computing in the Italian sites for the ALICE experiment at the LHC has been deployed. A combination of different complementary monitoring tools is typically used in most of the Tier-2 sites: this makes somewhat difficult to figure out at a glance the status of the site and to compare information extracted from different sources for debugging purposes. To overcome these limitations a dedicated ALICE dashboard has been designed and implemented in each of the ALICE Tier-2 sites in Italy: in particular, it provides a single, interactive and easily customizable graphical interface where heterogeneous data are presented. The dashboard is based on two main ingredients: an open source time-series database and a dashboard builder tool for visualizing time-series metrics. Various sensors, able to collect data from the multiple data sources, have been also written. A first version of a national computing dashboard has been implemented using a specific instance of the builder to gather data from all the local databases.

1. Introduction and motivation

The ALICE experiment at CERN was designed to study the properties of the strongly-interacting hot and dense matter created in heavy-ion collisions at the LHC energies [1]. The computing model of the experiment currently relies on the hierarchical Tier-based structure, with a top-level Grid site at CERN (Tier-0, also extended to the Wigner Datacenter in Budapest) and several worldwide-spread data centers at a regional level (Tier-1 and Tier-2 sites). The Italian computing infrastructure is mainly composed of a Tier-1 site at CNAF (Bologna) and four Tier-2 sites (Bari, Catania, Padova-Legnaro and Torino), with the additional contribution of two small Worldwide LHC Computing Grid (WLCG) centers in Cagliari and Trieste. Globally they contribute by about 15% to the overall ALICE computing resources [2].

The management of a Tier-2 site is based on a few complementary monitoring tools, each looking at the computing activity for ALICE from a different point of view: the ALICE-dedicated monitoring



system (MonALISA [3]) is used to extract information on the experiment side, the Local Batch System provides statistical data on the overall site activity and the Local Monitoring System provides the status of the computing machines. This scheme makes somewhat difficult to figure out at a glance the status of the site activity for the experiment and to compare information extracted from the different sources for debugging purposes. These limitations have triggered the design and development of an integrated monitoring system to be deployed in each of the ALICE Tier-2 sites in Italy: in the following sections details on the adopted technological solutions and reasons for evolving the local site dashboards towards a global national monitoring project are presented.

2. Basic technological solutions

A centralized site dashboard based on specific tools has been developed to meet tight technical requirements. Its design has been mainly guided by the need to manage large amounts of data in a fast way through an interactive and customizable Graphical User Interface. The present version of the dashboard, currently running in six of the ALICE INFN computing sites (Bari, Cagliari, Catania, Padova-Legnaro, Torino and Trieste), relies on InfluxDB [4], an open source time-series database able to reduce both the disk usage and the query time. Moreover, the query time is optimized thanks to the InfluxDB on-the-fly aggregation capability that allows one to send back to the client only the down-sampled series instead of a large amount of the full raw data.

The graphical component of the system is managed through Grafana [5], a dashboard builder tool for visualizing time-series metrics, which offers an easy interface to inspect and plot data. Specific sensors have been developed for each of the selected data sources: it allows one to extract the corresponding relevant information and then store it in the InfluxDB database, in such a way to take advantage from the specific aggregation capabilities mentioned above. The schematic view of the ALICE site dashboard architecture is shown in the figure 1.

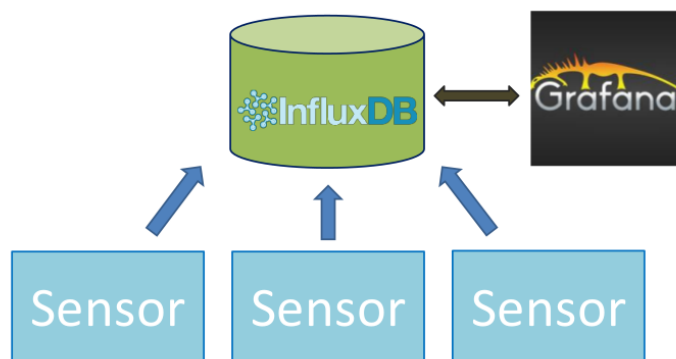


Figure 1. Schematic view of the ALICE site dashboard architecture.

3. Site dashboard for the ALICE Tier-2 in Bari

The first dashboard prototype has been deployed in the Bari ALICE Tier-2 site in 2014. The selected data sources for this datacenter are MonALISA, Zabbix [6] and HTCondor [7]. Zabbix is the local monitoring system by which all the machines are monitored and all the information related to the ALICE services are collected. HTCondor is a specialized workload management system for compute-intensive jobs: it provides data on the status of all the managed and terminated jobs belonging to all the queues, included the ALICE one. Dedicated sensors have been developed to extract relevant data from these sources and store them into the InfluxDB database. The Bari dashboard has been deployed

on a virtual machine provided from the local ReCaS [8] datacenter which hosts the ALICE Tier-2 activities: a screenshot of the ALICE dashboard for the ALICE Tier-2 site is shown in figure 2.




Figure 2. Example panel of the ALICE Dashboard for the Tier-2 site in Bari.

Despite the very limited resources required (2 VCPU, 4GB RAM and 10 GB Disk), the system provides good performance for the reactivity and long term data archiving have been achieved by such a system. Indeed, in almost two years running, more than 25 million measurements have been stored in InfluxDB using only 80 MB of disk space. With the currently allocated computing resources, the dashboard allows one to retrieve historical series covering two years in just few seconds.

4. National dashboard for the ALICE computing in Italy

The dashboard project initiated in Bari since 2014 has been further extended to the other Tier-2 sites and to the Trieste and Cagliari datacenters in 2016. For the Torino dashboard, an additional sensor needed to be developed in order to retrieve data from the local batch system (TORQUE [9]), while the sensor developed for the Bari site has been adapted to gather data from the local Zabbix. For the Padova-Legnaro site, new sensors was developed in order to allow data acquisition from the local batch system (LSF [10]) and for the local monitoring system (Nagios [11]). The preliminary systems currently running in Cagliari, Catania and Trieste at the moment collect data from MonALISA only. Having a dashboard running in all the selected Italian sites allowed us to start the implementation of a unique centralized monitoring system for the ALICE computing in Italy. Such a tool should be able to access the information collected in all the involved sites and monitor their performance: a screenshot of the preliminary version is shown in figure 3.

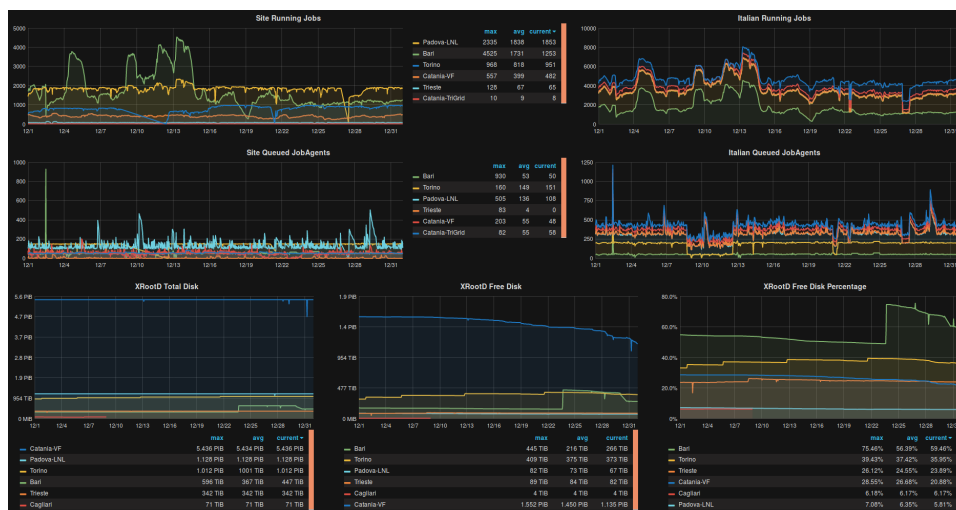


Figure 3. Glance of the preliminary version for the Italian ALICE Dashboard.

The national dashboard will also allow to compare parameters and perform a deeper investigation through the gathered information for debugging purpose and report making. All of this will be in addition supported by a simple and useful graphical interface. Its functionality will be more and more extended as additional information from the selected data sources will be collected through the local dashboard implementations in the involved sites.

5. Conclusion and outlook

A dashboard has been deployed in most of the Italian centers devoted to the computing for the ALICE experiment at the LHC. Such local implementations allow one to collect information from the different data sources present in the sites, store them in a dedicated InfluxDB database and present a detailed status of the site activity through a single and easily customizable Grafic User Interface based on Grafana. A preliminary version of an integrated national dashboard collecting information from the single InfluxDB instances has been also deployed and is currently under further development.

An evolution of the national dashboard project, aiming to make it more generic and with additional functionalities, is also in progress. The goal is ta system able to monitor cloud-federated datacenters based on Openstack. The key features of such a system will be the capability to organize a large amount of data, to evaluate the disruption cause roots and to prevent malfunctions before they happen. To implement these features, advanced machine learning and graph theory algorithms are being considered and tested.

Bibliography

- [1] ALICE Experiment, K. Aamodt et al., JINST 3 S08002 (2008).
- [2] WLCG, WLCG Technical Design Report, CERN-LHCC-2005-024 (2005).
- [3] MonALISA, <http://pcalimonitor.cern.ch/map.jsp>
- [4] InfluxDB, <https://docs.influxdata.com/>
- [5] Grafana, <http://grafana.org/>
- [6] Zabbix, <http://www.zabbix.com/>
- [7] HTCCondor, <https://research.cs.wisc.edu/htcondor/>
- [8] ReCaS Bari, <https://www.recas-bari.it/index.php/it/>
- [9] TORQUE, <http://www.adaptivecomputing.com/products/open-source/torque>
- [10] LSF, <https://portal.ictp.it/icts/ictpguide/lsf/>
- [11] Nagios, <https://www.nagios.org/>