CHEP 2016

San Francisco, October 10-14 2016

A Dashboard for the Italian Computing in ALICE

G. Vino[1,2], G. Donvito[2], D. Elia[2], A. Franco[2], S. Bagnasco[3], S. Piano[4]

[1] GARR, [2] INFN Bari – Italy, [3] INFN Torino – Italy, [4] INFN Trieste – Italy

on behalf of The ALICE Collaboration

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 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 WLCG centers in Cagliari and Trieste. Globally it contributes by about 15% to the overall ALICE computing resources [2].

Actually the management of a Tier-2 site is based on a few complementary monitoring tools, each looking at the ALICE activity from a different point of view: MonALISA [3] is used to extract information from 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 typical schema makes somewhat difficult to figure out at a glance the status of the ALICE activity in the site and to compare information extracted from different sources for debugging purpose. In this contribution, an integrated monitoring system for each of the ALICE Tier-2 sites and its evolution towards a national computing dashboard are presented.

Technological Solutions

A centralized site dashboard based on specific tools selected to meet tight technical requirements, like the capability to manage a huge amount of data in a fast way and through an interactive and customizable Graphical User Interface, has been developed. The present version of such dashboard, actually running in two of the ALICE INFN Tier-2 sites (Bari and Torino), 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 to return to the client only the down-sampled series instead of a large amount of 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 of information in order to extract corresponding relevant data from them and sequently to store them in InfluxDB in such a way to take advantage from database aggregation capabilities.

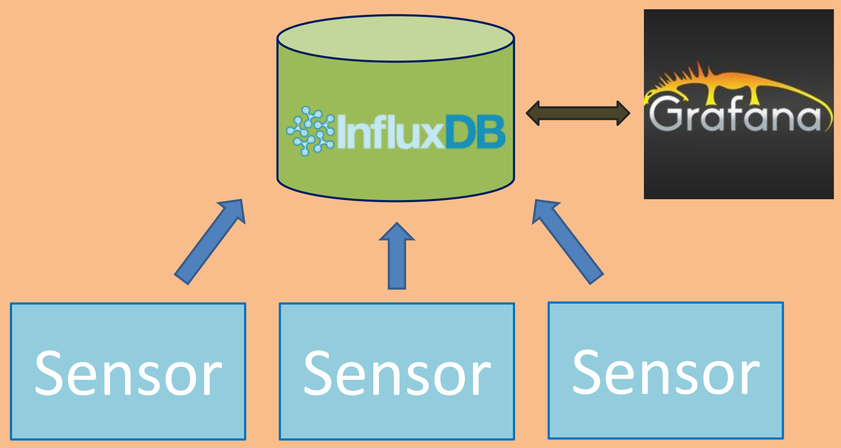

Bari Tier-2 Dashboard

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 and provides data on the status of all managed and terminated jobs belonging to all queues, included ALICE. Dedicated sensors have been developed to extract relevant data from these sources and to store them in InfluxDB.

The Bari dashboard has been deployed on a virtual machine provided from the local ReCaS [8] datacenter. Despite the very limited resources required (2 VCPU, 4GB RAM e 10 GB Disk), remarkable performances for the reactivity and long term data archiving have been achieved. 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 to retrieve historical series covering two years in just few seconds.

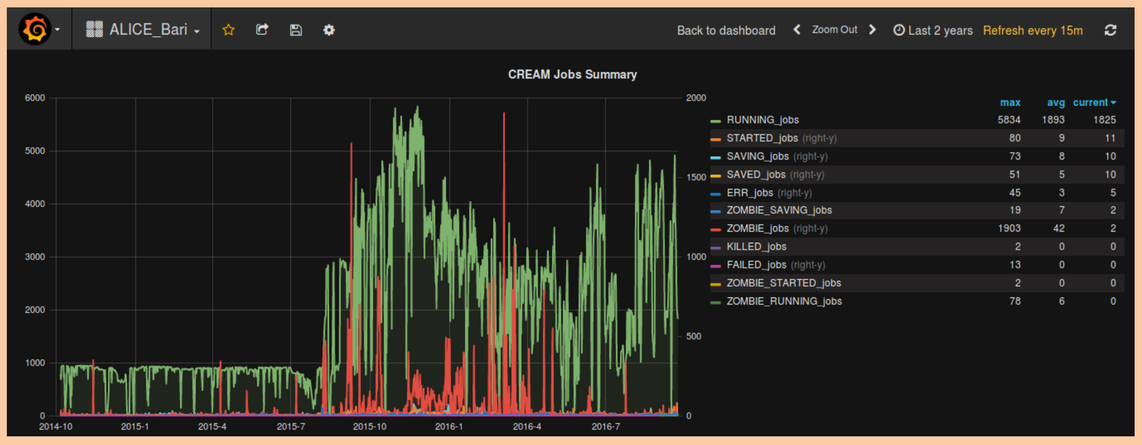

Italian Computing Dashboard

The dashboard project initiated in Bari has been extended also to the Torino ALICE Tier-2 site where a similar monitoring has been taking data for a few months already. For the Torino dashboard an additional sensor needed to be developed in order to retrieve data from the local Batch System (**PBS** [9]), while the already existing sensor has been adapted to gather data from the local Zabbix. The collected data are stored in the local InfluxDB. Also for Torino, the system has been deployed on a virtual machine with the same requirements proven to be adequate for the Bari dashboard. The project is currently being adapted also to the Padova-Legnaro ALICE Tier-2 site. New sensors are under development to allow acquiring data from the local data sources, namely **LSF** [10] and **Nagios** [11] as local batch system and monitoring system, respectively.

Following the currently running versions and the ongoing implementations, the dashboard project will be deployed also in the Trieste, Catania and Cagliari ALICE sites. Having a dashboard running in all the selected ALICE sites will allow the implementation of a unique centralized dashboard for the ALICE Tier-2 computing in Italy, a tool able to access the information collected in the sites and monitor their performance. This will also allow to compare parameters and investigate 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. Actually a preliminary version of such a national dashboard has been already implemented using only the data retrieved from the Bari and Torino sites. The functionality of such an Italian computing dashboard will be more and more extended as additional information from the selected sites will be collected through the local dashboard implementations.

An evolution of the current national dashboard project, aiming to make it more generic and with additional functionalities, is also in progress. The goal is the development of a system able to monitor cloud-federated datacenters based on Openstack. The key features of such 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.

References

[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] HTCondor, https://research.cs.wisc.edu/htcondor/
[8] ReCaS Bari, https://www.recas-bari.it/index.php/it/
[9] PBS, http://www.pbsworks.com/
[10] LSF, https://portal.ictp.it/icts/ictpguide/lsf/
[11] Nagios, https://www.nagios.org/