

A dashboard for the Italian Computing in ALICE

Tuesday 11 October 2016 16:30 (15 minutes)

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. 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 Wigner) and several globally distributed datacenters at national and regional level (Tier-1 and Tier-2 sites). The Italian computing infrastructure is mainly composed by a Tier-1 site at CNAF (Bologna) and four Tier-2 sites (at Bari, Catania, Padova-Legnaro and Torino), with the addition of two small WLCG centers in Cagliari and Trieste. Globally it contributes by about 15% to the overall ALICE computing resources.

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: for instance, MonALISA is used to extract information from the experiment side, the Local Batch System allows to store 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 purposes. In this contribution, a monitoring system able to gather information from all the available sources to improve the management of an ALICE Tier-2 site will be presented. 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 current version, running in the Bari Tier-2 site since more than one year, relies on an open source time-series database (InfluxDB): a dataset of about 20 M values is currently stored in 400 MB with on-the-fly aggregation allowing to return downsampled series with a factor of 10 gain in the retrieval time. A dashboard builder for visualizing time-series metrics (Grafana) has been identified as best suited option, while dedicated code has been written to implement the gathering phase. Details of the dashboard performance as observed along the last year will be also provided.

The system is currently being exported to all the other sites in order to allow a next step where a unique centralized dashboard for the ALICE computing in Italy will be implemented. Prospects of such an Italian dashboard and further developments on this side will be discussed. They also include the design of a more general monitoring system for distributed datacenters able to provide active support to site administrators in detecting critical events as well as in improving problem solving and debugging procedures.

Tertiary Keyword (Optional)

Cloud technologies

Secondary Keyword (Optional)

Accounting and information

Primary Keyword (Mandatory)

Monitoring

Author: ELIA, Domenico (INFN Bari)

Co-authors: FRANCO, Antonio (Universita e INFN, Bari (IT)); DONVITO, Giacinto (INFN-Bari); VINO, Gioacchino (Universita e INFN, Bari (IT)); BAGNASCO, Stefano (I.N.F.N. TORINO); PIANO, Stefano (INFN (IT)); SPINOSO, Vincenzo (Universita e INFN (IT))

Presenter: VINO, Gioacchino (Universita e INFN, Bari (IT))

Session Classification: Posters A / Break

Track Classification: Track 7: Middleware, Monitoring and Accounting