

Unified Monitoring Architecture for IT and Grid Services

Monday, 10 October 2016 14:45 (15 minutes)

For over a decade, LHC experiments have been relying on advanced and specialized WLCG dashboards for monitoring, visualizing and reporting the status and progress of the job execution, data management transfers and sites availability across the WLCG distributed grid resources.

In the recent years, in order to cope with the increase of volume and variety of the grid resources, the WLCG monitoring had started to evolve towards data analytics technologies such as ElasticSearch, Hadoop and Spark. Therefore, at the end of 2015, it was agreed to merge these WLCG monitoring services, resources and technologies with the internal CERN IT data centres monitoring services also based on the same solutions.

The overall mandate was to migrate, in concertation with representatives of the users of the LHC experiments, the WLCG monitoring to the same technologies used for the IT monitoring. It started by merging the two small IT and WLCG monitoring teams, in order to join forces to review, rethink and optimize the IT and WLCG monitoring and dashboards within a single common architecture, using the same technologies and workflows used by the CERN IT monitoring services.

This work, in early 2016, resulted in the definition and the development of a Unified Monitoring Architecture aiming at satisfying the requirements to collect, transport, store, search, process and visualize both IT and WLCG monitoring data. The newly-developed architecture, relying on state-of-the-art open source technologies and on open data formats, will provide solutions for visualization and reporting that can be extended or modified directly by the users according to their needs and their role. For instance it will be possible to create new dashboards for the shifters and new reports for the managers, or implement additional notifications and new data aggregations directly by the service managers, with the help of the monitoring support team but without any specific modification or development in the monitoring service.

This contribution provides an overview of the Unified Monitoring Architecture, currently based on technologies such as Flume, ElasticSearch, Hadoop, Spark, Kibana and Zeppelin, with insight and details on the lessons learned, and explaining the work done to monitor both the CERN IT data centres and the WLCG job, data transfers and sites and services.

Primary Keyword (Mandatory)

Monitoring

Secondary Keyword (Optional)

Tertiary Keyword (Optional)

Primary author: AIMAR, Alberto (CERN)

Co-authors: GARRIDO BEAR, Borja (Universidad de Oviedo (ES)); ZOLNAI, Daniel (Budapest University of Technology and Economics (HU)); KARAVAKIS, Edward (CERN); RIAHI, Hassen (CERN); RODRIGUEZ MARTINEZ, Javier (CERN); MAGNONI, Luca (CERN); GEORGIU, Maria-Varvara (Athens University of Economics and Business (GR)); SAIZ, Pablo (CERN); ANDRADE, Pedro (CERN); RAMA BALLESTEROS, Rocio (CERN); BELOV, Sergey (Joint Inst. for Nuclear Research (RU))

Presenter: KARAVAKIS, Edward (CERN)

Session Classification: Track 7: Middleware, Monitoring and Accounting

Track Classification: Track 7: Middleware, Monitoring and Accounting