

A Big Data Platform for heterogeneous data collection and analysis in large-scale data centers

Tuesday, 23 March 2021 13:40 (20 minutes)

With the upcoming start of Run-3, and especially Run-4, the amount of data managed by the WLCG data centres is expected to massively increase. In this context the importance of using efficient routines for the analysis of the incoming data is crucial to ensure the effectiveness of the HEP experiments. The actionable data that is available for the analysis also include the logging data produced by each computing node of the WLCG infrastructure.

Much information could be extracted from these sources, regarding the state of the system, and proper analysis of this data could result in better insight on its working state and eventually predict and prevent faults and errors. While efforts towards the extraction of these insight has enormous potential advantages, having an effective and efficient monitoring system is a critical asset, and a continuous concern for administrators of the data centre.

The INFN Tier-1 data centre hosted at CNAF has used various monitoring tools over the years all replaced, a few years ago, by a system common to all CNAF functional units, based on Sensu, Influxdb, and Grafana to acquire, store and visualize facility metrics (e.g. CPU load, memory usage, IO requests). Given the complexity of the inter-dependencies of the several services running at the data centre and the foreseen large increase of resource usage, a more powerful and versatile monitoring system is needed.

In this context, INFN-CNAF is promoting effort to introduce a log ingestion and analytics platform across its services, with the purpose of exploring possible solutions for the development of Predictive Maintenance model to detect and anticipate failures. This new monitoring system should provide a framework to provide and be able to correlate log files and metrics coming from heterogeneous sources and devices. Due to the novelty of this approach, and to the variety and heterogeneity of the involved data sources, identifying, extracting, and processing valuable information in an efficient and effective way is a challenging task.

We present a layered scalable big data infrastructure for data aggregation, exploration, and analysis for large data centers. Leveraging open source and de facto standard technologies, the proposed infrastructure follows a general approach ensuring its portability to different frameworks. On the bottom level, forwarding tools such as Filebeat, Fluentd and Thingsboard perform data ingestion from different sources (e.g. syslog, time series databases, IoT sensors). Timely distribution to processing nodes is carried out at a higher level by the topic-based publish-subscribe engine Apache Kafka. Processing is performed through Apache Spark and Spark Streaming instances. Data persistence is achieved through MinIO's high performance Object Storage, that allows long term storage for both raw data and the results of the analysis. The topmost layer includes visualization and alerting tools such as Kibana and Grafana through which the final users (i.e., sysadmins) may monitor the status of the system, be notified about predicted faults or explore new data.

The platform will satisfy the requirements of the Tier-1 group and the CNAF departments for more actionable insight from monitoring data, and provide the data center with an advanced infrastructure able to explore and analyze heterogeneous data in all the projects the center is involved.

Primary authors: Dr MARTELLI, Barbara (INFN - CNAF); CESINI, Daniele (INFN-CNAF); Mr ROSSI TISBENI, Simone (INFN-CNAF)

Co-authors: Dr FALABELLA, Antonio (INFN); CARBONE, Arianna (INFN-CNAF); CAVALLARO, Claudia (INFN-CNAF); MICHELOTTO, Diego (INFN-CNAF); DUMA, Doina Cristina (INFN - CNAF); FURLAN, Elisabetta (INFN-CNAF); Dr RONCHIERI, Elisabetta (INFN CNAF); SERGI, Giusy (INFN-CNAF); GASPARETTO, Jacopo (INFN-CNAF); MORGANTI, Lucia (INFN-CNAF); GALLETTI, Matteo (INFN-CNAF)

Presenter: Mr ROSSI TISBENI, Simone (INFN-CNAF)

Session Classification: Data Management & Big Data Session

Track Classification: Data Management & Big Data