EPJ Web of Conferences **295**, 01032 (2024) https://doi.org/10.1051/epjconf/202429501032
CHEP 2023

General purpose data streaming platform for log analysis, anomaly detection and security protection

Francesco Amori[1], *Stefano* Antonelli[1], *Vincenzo* Ciaschini[1], *Antonio* Falabella[1], *Enrico* Fattibene[1*], *Federico* Fornari[1], *Daniele* Lattanzio[1], *Diego* Michelotto[1], and *Lucia* Morganti[1]

[1]INFN-CNAF, viale Berti Pichat 6/2, 40127 Bologna, Italy

Abstract. INFN-CNAF is one of the Worldwide LHC Computing Grid (WLCG) Tier-1 data centres, providing computing, networking and storage resources to a wide variety of scientific collaborations, not limited to the four LHC (Large Hadron Collider) experiments. The INFN-CNAF data centre will move to a new location next year. At the same time, the requirements from our experiments and users are becoming increasingly challenging and new scientific communities have started or will soon start exploiting our resources. Currently, we are reengineering several services, in particular our monitoring infrastructure, in order to improve the day-by-day operations and to cope with the increasing complexity of the use cases and with the future expansion of the centre.
This scenario led us to implement a data streaming infrastructure designed to enable log analysis, anomaly detection, threat hunting, integrity monitoring and incident response. Such data streaming platform has been organised to manage different kinds of data coming from heterogeneous sources, to support multi-tenancy and to be scalable. Moreover, we will be able to provide an on demand end-to-end data streaming application to those users/communities requesting such kind of facility.
The infrastructure is based on the Apache Kafka platform, which provides streaming of events at large scale, with authorization and authentication configured at the topic level for ensuring data isolation and protection. Data can be consumed by different applications, such as those devoted to log analysis, which provide the capability to index large amounts of data and implement appropriate access policies to inspect and visualise information.
In this contribution we will present and motivate our technological choices for the definition of the infrastructure, we will describe its components and we will depict use cases which can be addressed with this platform.

1 Introduction

CNAF is the main technological centre of INFN (Italian Institute for Nuclear Physics). It hosts the main INFN data centre participating in the WLCG collaboration and supports the computing model of more than 30 scientific communities.

* Corresponding author: enrico.fattibene@cnaf.infn.it

With the aim of having a flexible service to manage large sets of heterogeneous data, such as monitoring/accounting information, service logs, reports of facility sensors, we deployed a modular data streaming infrastructure that can be used for troubleshooting and monitoring, for security protection (by implementing services for threat prevention, detection and response), for anomaly detection and for activities based on Machine Learning algorithms. Collected data can be consumed in different ways and analysed through specific tools or by developing custom utilities. Together with the monitoring services, this infrastructure allows Machine Learning studies or anomaly detection probes.

The paper is organised as follows. In Section 2 we describe the architecture and implementation of the platform, and in particular the log analysis infrastructure (Section 2.1) and the integration with a security monitoring solution (Section 2.2). Section 3 shows examples of how the infrastructure can be used to improve operations of the data centre. Finally, Section 4 closes the paper.

2 General purpose platform

The platform relies on open-source tools, suitably configured to offer multi-tenancy and complete segregation of data. Figure 1 shows the components of the infrastructure.

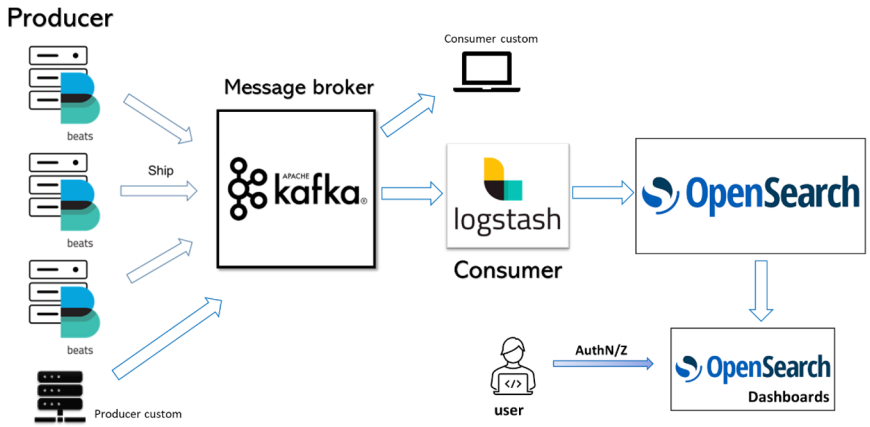

Fig. 1. Schema of the CNAF data streaming platform.

The core component is the message broker based on Apache Kafka [1], a distributed event streaming platform that implements message queuing. It can receive different types of data by different kinds of producers, such as those of the Beats [2] family, each of them devoted to ship data of different kinds (metrics, network data, logs, etc.). One of these tools is Filebeat, specialised in collecting and sending log file entries. It is possible to use custom utilities able to communicate with Kafka as data producers.

Kafka organises data in topics, logical spaces where information is put waiting to be consumed by predefined tools or custom scripts. By exploiting the Kafka ACL system, different CNAF departments or projects can use only their own topics, that is complete data segregation is guaranteed.

Currently Kafka is implemented as a three-node cluster composed of virtual machines hosted by three different virtualization systems, in order to ensure the availability of the service. Data is replicated to minimise the risk of loss; the administrators can decide the proper number of replicas for each topic.

2.1 Log analysis infrastructure

As a first use case of data consumers for Kafka, we deployed a log analysis infrastructure running a set of open-source tools derived from the ELK (Elasticsearch, Logstash, Kibana) stack.

Logstash [3] can parse data consumed from Kafka, as well as enrich them by creating fields based on information present in the log records. Data managed by Logstash can be sent to a variety of outputs, including Kafka itself to offer transformed data to its possible consumers.

In our system, data filtered by Logstash are sent to OpenSearch [4], a software derived from Elasticsearch and designed to search and analyse large amounts of data, organised in indexes. We chose OpenSearch to have a complete multi-tenancy service, thanks to its Security plugin for authentication and access control, which allows the administrator to manage users, tenants, permissions and roles in a flexible way. Additional plugins help users to create Machine Learning models or implement anomaly detection triggers.

As shown in Figure 2, the OpenSearch service is composed of three virtual nodes that run the Master role and four physical Data nodes that store indexes on SSD disks.

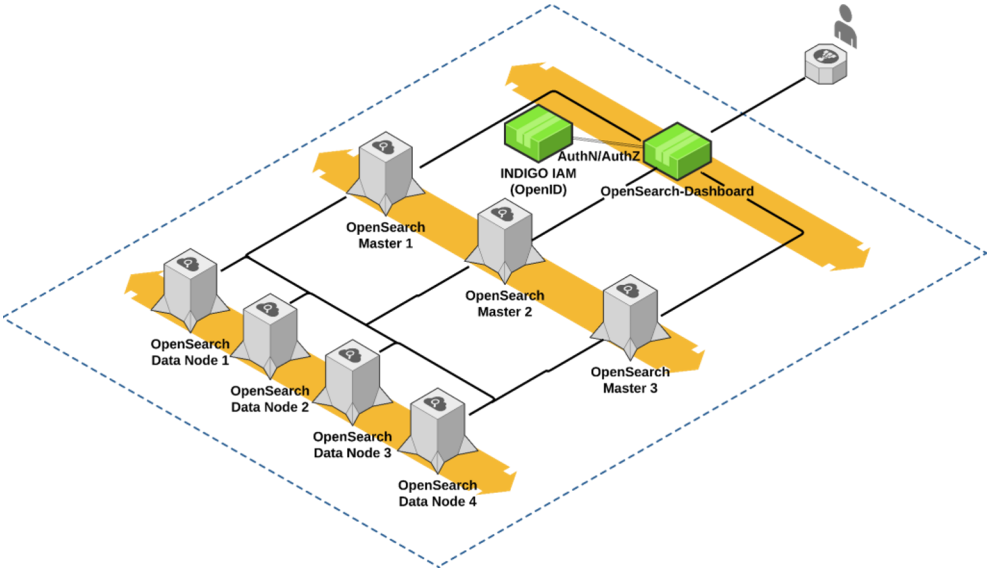

Fig. 2. OpenSearch service schema.

Data can be accessed through OpenSearch Dashboard or Rest APIs. In both cases, users can authenticate to OpenSearch through the INDIGO Identity and Access Management (IAM) service [5] that supports local users as well as INFN Authentication and Authorization Infrastructure. Local authentication through username and password is also supported, but used only for administrators and service users.

Authorization is managed thanks to the automatic mapping between the groups defined in IAM and specific roles defined within OpenSearch, configured by the administrators. In this way, users can have complete access to data of their department/project, and read-only capabilities on other data, depending on the type of analysis they need to perform. Multi-tenancy allows users of the same department/project to access common visualisations or dashboards inside their own tenant.

Similarly to Kafka, data are replicated to minimise the risk of loss; the administrators can choose the number of replicas of the indexes. OpenSearch can also create snapshots of groups of indexes using an external storage space. In our implementation, we chose a Ceph-based filesystem where snapshots are created on a regular basis. Since OpenSearch does not have a native quota management, we deployed an alarm system alerting the OpenSearch administrators in case the value of used space for each project goes beyond a configurable limit.

2.2 Integration with Wazuh

In order to improve CNAF security monitoring, we decided to integrate Wazuh [6] with OpenSearch. Wazuh is a well-known ossec-based and elastic-based software focused on security monitoring and reaction. We consider it the best solution in this field among freeware choices. Going beyond traditional intrusion detection, eXtended Detection and Response (XDR) proactively identifies and mitigates evolving cyber threats. The Host-Based Intrusion Detection System (HIDS) serves as a vigilant sentinel, swiftly alerting to potential breaches. Integration with OSSEC Active Response automates threat responses, reducing response times. Wazuh requires the installation of an agent on each monitored host and it mainly works by analysing the information sent back by the agents (chiefly, but not only, the log files) looking for behavioural patterns that correspond to dangerous or anomalous actions, and its algorithm may be expanded via integration with third party sources like a MISP (Malware Information Sharing Platform).

An instance of Wazuh has been set up using the platform's own OpenSearch as backend for the data and with a custom dashboard provided by Wazuh itself, with the idea to eventually monitor all CNAF nodes. Figure 3 shows the integration of the Wazuh services with OpenSearch.

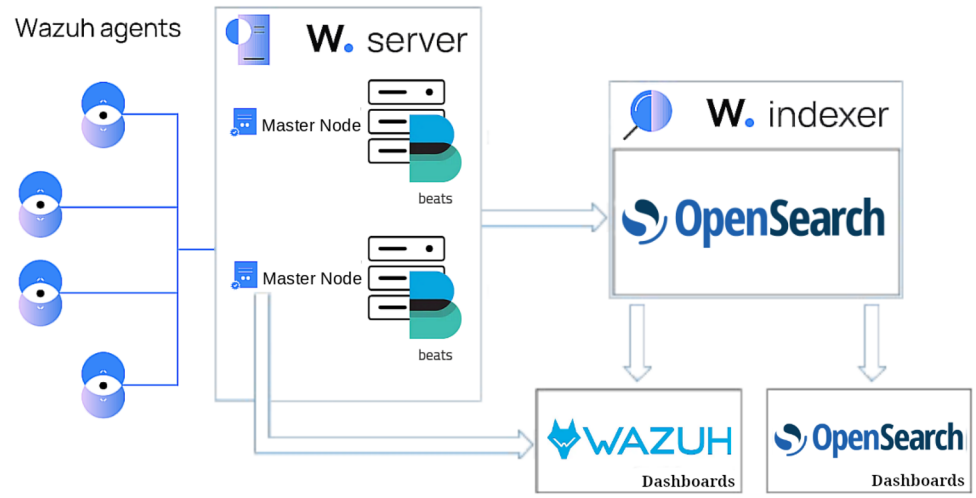

Fig. 3. Wazuh integration with OpenSearch.

Our current implementation is currently composed of two Master nodes which receive data from the Agents and, after some local re-elaboration, send it to the platform's OpenSearch service (called indexer by Wazuh) to be indexed and easily retrieved. There is also a dashboard, for the moment independent from the platform's main one, which exposes the analysis, alerting and remediation interface. For the future, we intend to integrate our

instance with external sources of information like MISP, and experiment with the OpenSearch ML engine to apply Machine Learning techniques to the data as an additional analysis mechanism.

3 Use case: data centre operations

The monitoring system currently in place at INFN-CNAF data centre [7], which is based on Sensu, InfluxDB and Grafana, allows to inspect and visualise plenty of numerical information in its time series, and such information is used on a daily basis for checking and debugging the status of resources and services provided to the experiments.

Such information can be valuably enriched with highly specific metadata that can be obtained through the parsing and processing of the log files of services running at the data centre through the log analysis infrastructure outlined in the above sections.

This is of great help when it comes to quick detection of problematic patterns or irregularities and to thorough debugging of the status of services in support of day-by-day troubleshooting and operations.

Beyond log files coming from other services, we are currently exploiting the log analysis platform for all the data transfer and data management services running at INFN-CNAF; approximately 60 servers running this kind of services ingest their log files to the platform in a continuous live-feed streaming.

Each record is parsed via appropriate grok filters based on regular expressions in Logstash, to generate structured data which can then be searched and visualised within OpenSearch.

Figures 4, 5 and 6 provide different examples of the kind of data visualisations and insights which can be obtained from the platform. For instance, Figure 4 allows to compare the number of files transferred with different protocols (http/gsiftp): this kind of plot represents an essential tool in supporting and monitoring the transition from gsiftp to http which is currently ongoing within the Worldwide LHC Computing Grid. The number of files transferred via StoRM WebDAV [8] using https protocol or via GridFTP using gsiftp protocol can only be inferred from the log files of the services, counting successful transfers therein.

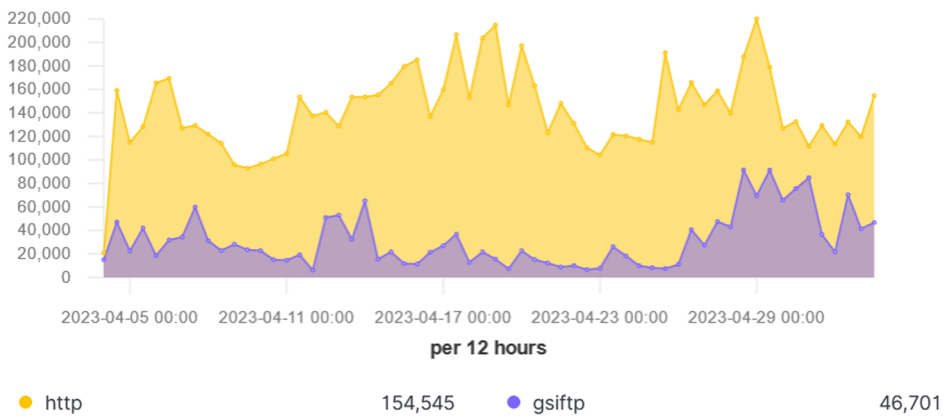

Fig. 4. Number of files transferred with different protocols, monitoring the ongoing transition from gsiftp to https in WLCG.

Instead, Figure 5 shows, for the specific ATLAS experiment, the number of http Third-Party-Copies, i.e. bulk data transfers between sites, together with their throughput. This is yet another information which can only be inferred from entries in the log files and which can help to quantify the load and stress of the servers involved in the data transfers and to identify possible slowing down of the services.

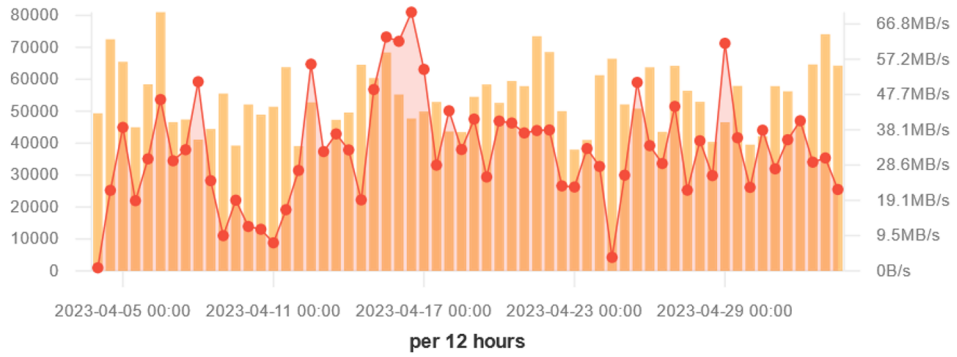

Fig. 5. Number and throughput of http Third-Party-Copies for the ATLAS experiment.

A further example is provided with Figure 6, where the platform allows to plot the distribution of the requests of file transfers among the four different StoRM WebDAV endpoints dedicated to the ATLAS experiment. This represents an immediately insightful visualisation to spot possible load balancing/unbalancing of requests.

In general, much valuable information is enclosed in the log files of these services and can be inspected and visualised with this platform: the number and type of requests, their distribution on the servers, the geographical origin of the requests, the storage areas on which such requests insist, the rate of success and failures, the kind and frequency of requests or errors, the biggest players in case of shared resources, and many more.

Thus, the great outcome of this activity is the availability of a useful interface for CNAF system administrators, providing both a powerful search engine to correlate many different log files and many tools for interactive data discovery and visualisations in support of their everyday debugging activity and operations.

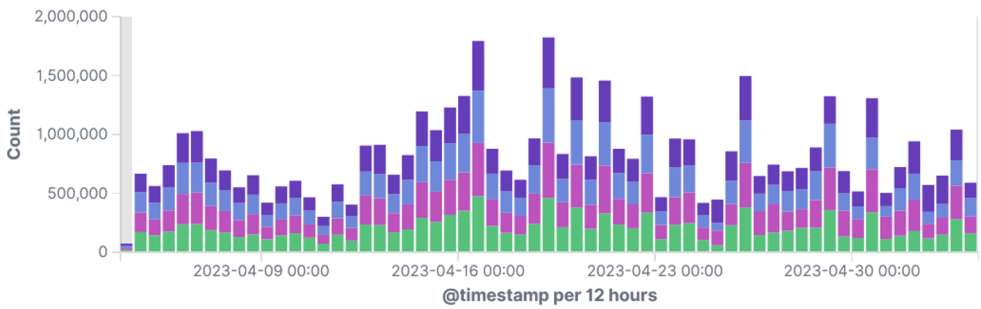

Fig. 6. Distribution of the requests of file transfers among the four different StoRM WebDAV endpoints of the ATLAS experiment at INFN-CNAF data centre.

4 Conclusions and future works

In this contribution we described a flexible platform to manage large sets of heterogeneous data, such as monitoring/accounting information, service logs, reports of facility sensors. The platform is in production since Spring 2023 and receives log data from tens of servers running at INFN-CNAF data centre. The implementation of the services has been made in order to provide high availability, scalability, multi-tenancy and data segregation. Moreover, the platform is integrated with other services running at CNAF, such as the provisioning infrastructure to automatically configure the nodes and the monitoring system to check the status of the components.

This service allows operators to have flexible tools to enhance their troubleshooting capabilities. Moreover, data can be analysed by administrators or users for a variety of use cases, also with the help of Machine Learning techniques. In addition, the integration of Wazuh with OpenSearch enhances the capabilities of CNAF security monitoring and, at the same time, minimises the operations effort, by exploiting an already running service.

For the future, we plan to investigate and test OpenSearch additional plugins, such as Anomaly Detection to build specific sensors directly applicable to data centre operations, and Machine Learning to train neural networks for failure prediction.

References

1. Apache Kafka. https://kafka.apache.org. Accessed 6 Dec. 2023
2. Beats. https://www.elastic.co/beats. Accessed 6 Dec. 2023
3. Logstash. https://www.elastic.co/logstash. Accessed 6 Dec. 2023
4. OpenSearch. https://opensearch.org. Accessed 6 Dec. 2023
5. A. Ceccanti, E. Vianello, M. Caberletti, F. Giacomini, *Beyond X.509: token-based authentication and authorization for HEP,* in Proceedings of 23rd International Conference on Computing in High Energy & Nuclear Physics, CHEP 2018, 9-13 July 2018, Sofia, Bulgaria (2018)
6. Wazuh. https://wazuh.com. Accessed 6 Dec. 2023
7. S. Bovina, D. Michelotto, *The evolution of monitoring system: the INFN-CNAF case study*, 2017 J. Phys.: Conf. Ser. 898 092029
8. A. Ceccanti, E. Vianello, F. Giacomini, *Token-based authorization in the StoRM WebDAV service*. 24th International Conference on Computing in High Energy & Nuclear Physics, CHEP 2019, 4-8 November 2019, Adelaide, Australia (2019)