Contribution ID: **151** Type: **Oral Presentation**

A Big Data Platform for Scalable Monitoring and Analytics at INFN-CNAF

Thursday, 19 March 2026 17:00 (30 minutes)

An heterogeneous data management and analytics platform has been deployed over the years at INFN-CNAF in Bologna, with the goal to support both system administrators and user-support teams in monitoring, diagnostics, and operational decision-making across the data center, that serves multiple particle and astroparticle experiments (and beyond) in Italy and internationally. The platform is built around a general-purpose message-handling backbone based on Apache Kafka, enabling scalable ingestion of high-volume, high-velocity data streams. On top of this backbone, a log-analysis stack - comprising Logstash, OpenSearch, and OpenSearch Dashboards, the open-source fork of the ELK suite maintained by AWS - provides powerful indexing, search, and visualization capabilities, including advanced authentication and authorization plugins. This contribution presents the architecture of the platform and the design choices that guided its implementation, complemented by a use case on farming services that demonstrates the platform's applicability in real-world scenarios . We also report on the extensions and enhancements introduced to transform the platform into a robust environment for Big Data analytics and AI applications aimed at improving data-center monitoring, alerting, and operational control. Finally, we illustrate the range of user-facing capabilities that make the platform a flexible and accessible tool for advanced analytics.

Primary authors: MICHELOTTO, Diego (INFN-CNAF); LORUSSO, Marco (University of Bologna and INFN)

Presenter: LORUSSO, Marco (University of Bologna and INFN)

Session Classification: Infrastructure Clouds and Virtualisations - III

Track Classification: Track 8: Infrastructure Clouds and Virtualizations