

Big data analytics as a service infrastructure: challenges, desired properties and solutions

Manuel Martín-Márquez

CERN openlab, Information Technology Department, Geneva, Switzerland

E-mail: Manuel.Martin.Marquez@cern.ch

Abstract. CERN's accelerator complex generates a very large amount of data. A large volumen of heterogeneous data is constantly generated from control equipment and monitoring agents. These data must be stored and analysed. Over the decades, CERN's researching and engineering teams have applied different approaches, techniques and technologies for this purpose. This situation has minimised the necessary collaboration and, more relevantly, the cross data analytics over different domains. These two factors are essential to unlock hidden insights and correlations between the underlying processes, which enable better and more efficient daily-based accelerator operations and more informed decisions. The proposed Big Data Analytics as a Service Infrastructure aims to: (1) integrate the existing developments; (2) centralise and standardise the complex data analytics needs for CERN's research and engineering community; (3) deliver real-time, batch data analytics and information discovery capabilities; and (4) provide transparent access and Extract, Transform and Load (ETL), mechanisms to the various and mission-critical existing data repositories. This paper presents the desired objectives and properties resulting from the analysis of CERN's data analytics requirements; the main challenges: technological, collaborative and educational and; potential solutions.

1. Introduction

CERN's particle accelerator infrastructure is complex and heterogeneous. Several mission-critical subsystems, which represent cutting-edge technology in several engineering fields, are involved. Cryogenics, power converters, magnet protection and vacuum are among them. Over the last decades, aiming to ensure efficient and safe operations for each of these systems and therefore for the complete accelerator complex, CERN has implemented a complete infrastructure of custom data-driven control and monitoring systems. Several millions of control devices (including sensors, front-end equipment), constitute these critical-mission services and have led to a significant investment in data storage and management.

Fully exploiting these data and evolving the controls and monitoring infrastructures to intelligent, predictive and proactive control and monitoring systems become an essential challenge. Overcoming this will help us gain a deeper understanding of the factors that affect the performance and availability of the accelerator complex. This is critical for achieving maximum operational efficiency by ensuring more informed decision-making. CERN's Beams Control [3], Engineering [8] and Information Technology [6] departments have documented this fact.

With this objective in mind, CERN's equipment experts, operators and engineering teams have applied different data persistency and analytics approaches, techniques and technologies.



This situation has minimised the necessary collaboration and, more relevantly, the cross data analytics over different domains.

Together with the aforementioned facts, these two important factors have triggered the need for research and development activities regarding two major aspects. The first concerns the design based on real CERN analytics use cases. The second involves an implementation of a Big Data Analytics infrastructure as a Service (DAaaS), which aims to: (1) integrate the existing developments; (2) centralise and standardise the complex data analytics needs for the CERN's research and engineering community; (3) deliver real time, batch data analytics and information discovery capabilities; and (4) Provide transparent access and ETL, mechanisms to the different and mission-critical existing data repositories.

In addition, in spring 2015, the second Run of the Large Hadron Collider (LHC) was initiated, following a two-year period of maintenance and upgrade work. During the second run, the LHC will significantly increase its luminosity, and the energy per beam will almost double (Table 1) [4]. This is an exciting moment for the scientific community, but this also leads to important challenges for all of the subsystems that comprise CERN's accelerator complex and the various engineering teams responsible for them. This makes detailed analysis of the control and monitoring data, and the implementation of the DAaaS even more important.

Table 1. An overview of performance-related parameters during LHC operations in 2010, 2013 and design values.

Parameter	2010	2011	2012	design values
Beam Energy (TeV)	3.5	3.5	4.0	7.0
β^* in IP 1 and 5 (m)	2.0/3.5	1.5/1.0	0.6	0.55
Bunch spacing (ns)	150	75/50	50	25
Max. number of bunches	368	1380	1380	2808
Max. bunch intensity (protons per bunch)	1.2×10^{11}	1.45×10^{11}	1.7×10^{11}	1.15×10^{11}
Normalised emittance at start of fill	≈ 2.0	≈ 2.4	≈ 2.5	3.75
Peak luminosity $cm^{-2}s^{-1}$	2.1×10^{32}	3.7×10^{33}	7.7×10^{33}	1×10^{34}
Max. number of events per bunch crossing	4	17	37	19
Stored beam energy (MJ)	≈ 28	≈ 110	≈ 140	362

2. Challenges and desired features

2.1. Data repositories

As previously mentioned, CERN has very successfully developed custom data-driven control and monitoring systems. As a result of those systems, different data repositories with different objectives and domains have been deployed some relevant examples are:

- Control Configuration Database (CCDB) [11]: CCDB maintains data regarding the configuration required for the correct functioning of CERN's accelerators controls system. It covers the need for the configuration of components of the Controls System itself, e.g. the Controls Middleware (CMW), as well as accelerator components as seen by the Controls System, e.g. power converters, for all accelerators: the Large Hadron Collider (LHC), the Super Proton Synchrotron (SPS) Complex, the Proton Synchrotron (PS) Complex, and the CLIC Test Facility (CTF3).
- CERN Accelerator Logging Service (CALS) [2, 7]: CALS (Fig. 1) is used to store and retrieve billions of control data acquisitions per day, from across the complete CERN accelerator complex, related subsystems, and experiments.

- Accelerator Fault Tracking (AFT): AFT provides the infrastructure necessary to consistently and coherently capture, persist and make available accelerator fault data for further analysis.

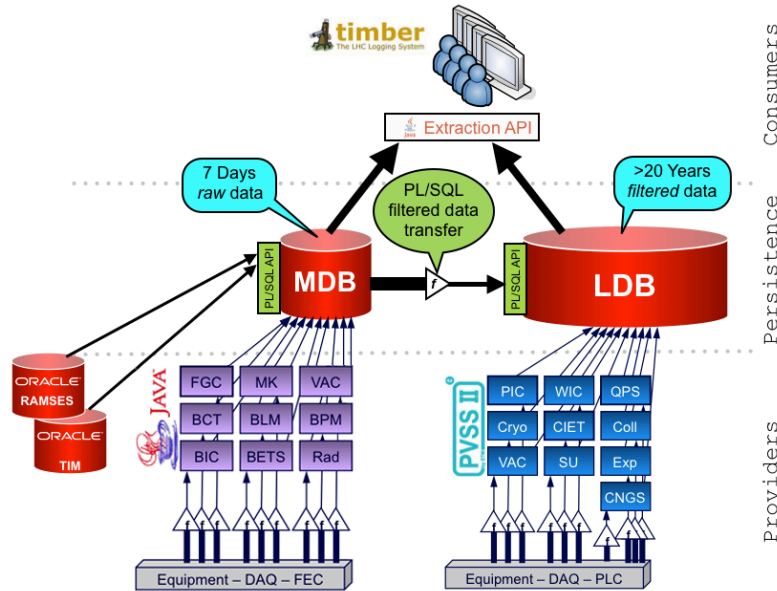


Figure 1. Logging Service Architecture Overview

Those various data repositories contain important information that needs to be transparently integrated in order to enable data analysis between different areas. Therefore, due to the extreme heterogeneity of the data sources, the ETL processes, required for loading the data from the repositories while standardising and reshaping it to make it suitable for complex analyses. Analyses are challenging, given the variety of formats and the nature of the data gathered and stored. In addition to the integration with the existing data repositories, the pending requirements result in the need to evaluate a potential general solution, capable of:

- storing huge amounts (hundreds of TBs) of structured and unstructured data safely; and for the long term.
- enabling flexible and varied data access technologies.

2.2. Near-real-time processing

The challenges of near-real-time processing focus on the ability to process large volumes of data (GBs per second) with low latency (on the order of seconds) originating from different sources and domains. The tools and techniques should be sufficiently flexible to import and apply the knowledge inferred from the data (batch analysis), or the pre-existent (human) knowledge of the systems. They may include models, pattern definitions and matching, thresholds and, most importantly, the ability to trigger actions based on the discovery of complex events within the streaming data. Due to the critical nature of the services the different systems offer, and the huge data volumes they produce and manage, it is important to consider some mandatory aspects, such as scalability, fault-tolerance, and the ability to guarantee that all the streaming data produced are processed and analysed. In addition to the aforementioned factors, the near-real-time processing system needs to integrate the different domains and analysis technologies that currently exist at CERN. Therefore, the support of a wide range of data analysis tools and programming languages is essential.

2.3. Batch processing

Analysing data originating from CERN control data repositories requires learning from the past and applying that knowledge to better understand current status and to predict future behaviour. This is the objective of batch analytics. Thus, the main challenge in this area is to enable the capability to mine and analyse huge volumes of structured and unstructured data coming from various repositories.

Most near-real-time analysis methods that exist today are based on our knowledge about the underlying systems. However, there are patterns, correlations and predictions among large quantities of data that cannot be immediately captured by the experts. These could be discovered by other means, like machine learning methods. This would provide much more reliable real-time or predictive information. Accordingly, one of the main challenges for both the present and the future. Linear and non-linear modelling, classical statistical tests, complex time series analysis and forecasting, classification, and clustering are just some of many techniques used with this purpose. An expected consequence of this type of analysis is the possibility of improving not only the measurable efficiency of the systems, but also the relative users' perception. This entails deep analysis of all applications, including study of their usage through log files and other means. It is envisaged that such predictive analytics techniques can not only improve the efficiency of the systems, but also help identify areas in which improvements or investments should be focused.

In addition, other important factors to consider, due to the volume of data to be handled and diverse nature of the analyses to be developed, are:

- allow spreading the data computation to multiple nodes, while performing the analysis as close to the data as possible.
- enable data access to multiple analytics frameworks and tools like R, Python, Spark (MLlib, SQL) and Cloudera Impala, empowering the capabilities to easily prototype potential models.

2.4. Information discovery

Information discovery is about much more than just representing data; it also allows us to unlock understanding, facilitate exploration, and identify hidden insights, correlations and patterns. Therefore, any information discovery solution must enable rapid integration of different data sources and exploitation of the data in a simple, efficient and self-service manner. These facts are consistent with more intelligent decision-making processes for improving the efficiency of operations at all levels. Fig. 2 shows an example of information discovery tool used in the context. In our domain, we have to consider two relevant factors:

- volume and diversity of users and areas of expertise involved.
- speed of implementation: prototyping, validating assumptions, and delivering results as quickly as possible from different perspectives are necessities in our environment.

3. Training and educational aspects required for the projects

The training and educational aspect in Data Analytics is a particularly important and challenging area. Many domains of expertise are involved and therefore need to be covered. These domains include the following: (1) analytics lifecycle: roles definitions, analytics phases; (2) basic methodologies and techniques: exploratory and statistical models; (3) advanced models: classification, clustering, association rules, decision trees, linear and logistic regressions, time series analysis and forecasting, text analysis; (4) high-level solutions: NoSQL ecosystem, complex event processing, in-database analytics, advanced SQL techniques, distributed near real-time computation system. Also regarding training, we have to consider the strong relationship with the previous point described in this paper.

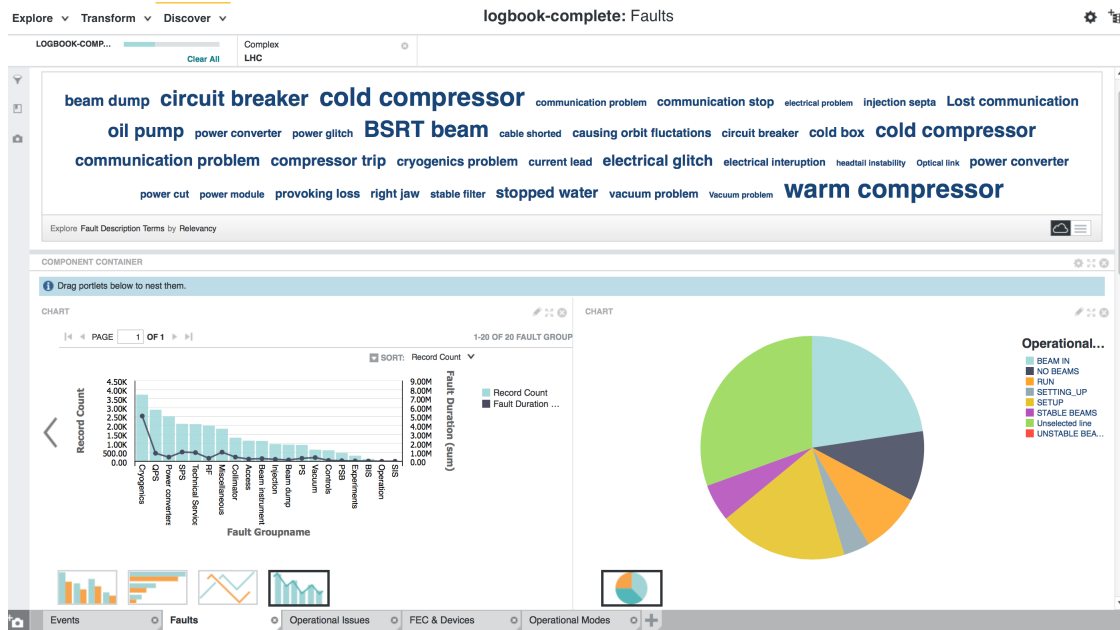


Figure 2. Big Data Information Discovery analysis interface: analysis of LHC faults extracted from logbook database

4. Big Data Analytics as a Service

4.1. Integration

The features defined in Section 2 need to be grouped to provide a self-contained solution for the Big Data Analytics as a Service infrastructure proposed in this paper (Fig. 3). The platform is defined as a set of layers in which different technologies coexist, since it is clear from the requirements that there is no one-size-fits-all solution. This set is encapsulated in a common framework that is capable of easily transferring data between the layers and the tools, so that the analyses can be performed with the most appropriate solutions. An important factor considered during the whole design process is that data analytics platform should be accessible and adaptable for some of the use cases to external institutions that are willing to cooperate. It should, therefore, use open and well-defined standards for exchanging the data. This is also important in terms of future support.

4.2. Layers

4.2.1. Data integration layer - The data acquisition layer is used as the data entry point. This layer takes data from the different repositories - structured and unstructured - and loads it into a common central distributed repository, based on the Hadoop ecosystem. This layer also provides some data transformation capabilities for those cases in which the type of analysis, and therefore the final data structure, are known beforehand. In this layer, two technologies with different scopes are applied:

- Apache Sqoop: This technology is used to transfer efficiently data from relational databases to an Apache Hadoop storage ecosystem.
- Flume: This Apache project collects and processes log data into our Hadoop-based storage layer and provides on-line aggregation capabilities.

4.2.2. Data storage layer - The data storage layer must enable efficient storage of large volumes of heterogeneous data and offer flexible and independent data access to the upper analytics layer. The technologies deployed in this layer are:

- Apache HDFS [1]: Hadoop Distributed Filesystem, which is used as the principal storage system.
- Apache HBASE [9]: Hadoop database, which provides distributed, scalable and big data store.
- TACHYON [5]: A memory-centric distributed storage system. Currently, we are evaluating the potential benefits of this technology, relative to Spark and MapReduce performance.

4.2.3. Data analytics layer - This is the front-end layer with which the analytics user will interact. From the very early stages of our design, it was clear that this layer needs to offer different solutions for integrating the existing analytics development and tools ranging from lower level tools like Spark [10], Python or R (including Oracle R Enterprise, and Oracle R connectors for Hadoop) to higher levels such as Impala, iPython, Zeppelin and self-service information discovery platforms such as Oracle Big Data Discovery.

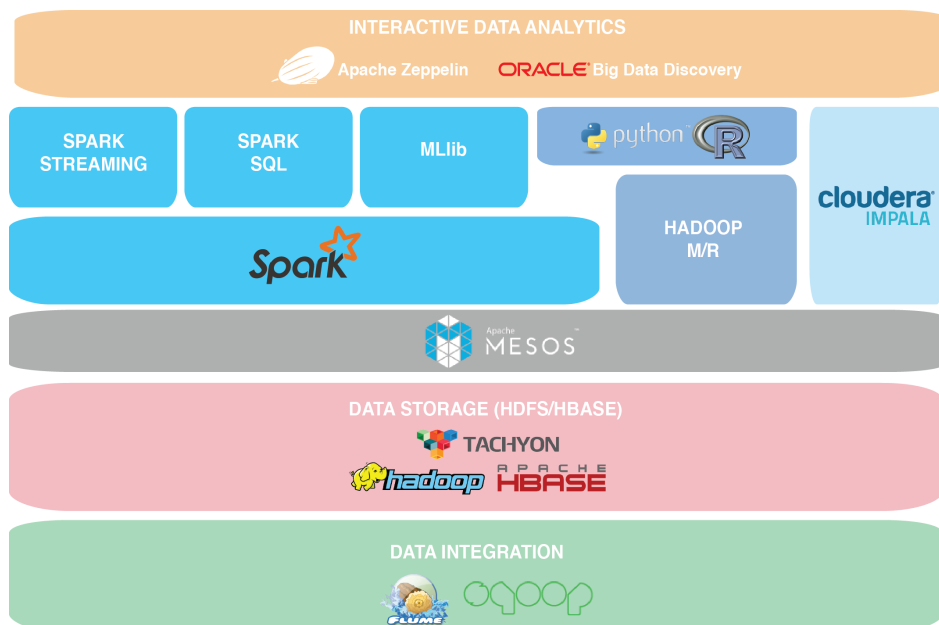


Figure 3. Modular Big Data Analytics as a Service Components

5. Conclusions

The control and monitoring of CERN's accelerators are extremely complex. Over the last decades, CERN has dedicated significant efforts to developing those systems, which has led to important data investments. Evolving those systems to be predictive, proactive and more intelligent has become essential to improving daily operations at all levels. The proposed Big Data Analytics as a Service Infrastructure described in this paper is the result of the requirements analysis of a significant number of use cases collected. It offers a centralised point to exploit those data to gain better understanding about the past, present and future behaviour of the factors driving the accelerator complex performance, availability and operations. The infrastructure also fulfils the main requirements: (1) data integration of the diverse existing repositories and future

developments, providing transparent access and ingestion mechanisms; and (2) data analytics requirements, enabling real-time, batch data analytics and information discovery capabilities. Nevertheless, the infrastructure is open and in continuous evolution due to the rapid pace the technology is following in the area of Big Data and analytics. This fact, and its potential negative implication, are mitigated by the modular design implemented, in which different technologies can coexist and easily be replaced in the future without affecting the users, the integrity of the data, or the overall infrastructure.

Acknowledgements

The author would like to thank Oracle and CERN openlab framework for supporting the work described in this paper. I also gratefully thank Chris Roderick (Beams department) who provided insight and expertise that greatly assisted on data management practices in general and on the Accelerator Fault Tracking and CERN Central Logging Service systems in particular. Thanks to Filippo Maria Tilaro, Axel Voitier and Benajin Bradu (Engineering department) for their inputs on the control systems data analytics requirements and use cases; and Antonio Romero Marin (CERN openlab) who helped with the technical stuff. And finally the author is also immensely grateful to Eric Grancher and Eva Dafonte Perez (Information Technology department) for their pearls of wisdom on the subject, support on the research, and comments on an earlier version of the manuscript.

References

- [1] Dhruva Borthakur. Hdfs architecture guide. *Hadoop Apache Project*, page 53, 2008.
- [2] G. Kruk C. Roderick, L. Burdzanowski. The cern accelerator logging service -10 years in operations: a look at the past, present, and future. *Proceedings of ICALEPCS, San Francisco, CA, USA*, 2013.
- [3] Steen Jesen. “what you get” – controls software. *LHC Beam Operation workshop - Evian*, 2012.
- [4] M. Lamont. High-quality beam from the injectors and full exploitation of options in the collider underpinned the lhc’s performance in 2010–2013. *CERN Courier*, 2013.
- [5] Haoyuan Li, Ali Ghodsi, Matei Zaharia, Scott Shenker, and Ion Stoica. Tachyon: Reliable, memory speed storage for cluster computing frameworks. In *Proceedings of the ACM Symposium on Cloud Computing*, pages 1–15. ACM, 2014.
- [6] Manuel Martin-Marquez. Big data analytics for improving the cern’s large hadron collider operations. *Big Data and Innovation Summit*, 2013.
- [7] C Roderick. Cern accelerator data logging and analysis. In *Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC), 2013 IEEE*, pages 1–3. IEEE, 2013.
- [8] Axel Voitier. Industrial controls data analytics use cases. *Openlab workshop on Data Analytics*, 2012.
- [9] Mehul Nalin Vora. Hadoop-hbase for large-scale data. In *Computer Science and Network Technology (ICCSNT), 2011 International Conference on*, volume 1, pages 601–605. IEEE, 2011.
- [10] Matei Zaharia, Mosharaf Chowdhury, Michael J Franklin, Scott Shenker, and Ion Stoica. Spark: cluster computing with working sets. In *Proceedings of the 2nd USENIX conference on Hot topics in cloud computing*, pages 10–10, 2010.
- [11] Zornitsa Zaharieva, M Peryt, and M Martin Marquez. Database foundation for the configuration management of the cern accelerator controls systems. In *Conf. Proc.*, number CERN-ATS-2011-206, page MOMAU004, 2011.