

# A PROTOTYPE FOR THE EVOLUTION OF ATLAS EVENTINDEX BASED ON APACHE KUDU STORAGE

Z. Baranowski<sup>1</sup>, L. Canali<sup>1</sup>, A. Fernandez Casani<sup>2</sup>, E. Gallas<sup>3</sup>, C. Garcia Montoro<sup>2</sup>, S. Gonzalez de la Hoz<sup>2</sup>, J. Hrivnac<sup>4</sup>, F. Prokoshin<sup>5</sup>, G. Rybkin<sup>4</sup>, J. Salt<sup>2</sup>, J. Sanchez<sup>2</sup>, D. Barberis<sup>6</sup>  
on behalf of the ATLAS Collaboration

<sup>1</sup>CERN, Geneva, Switzerland; <sup>2</sup>Institut de Física Corpuscular, Valencia Spain; <sup>3</sup>University of Oxford, Oxford, UK; <sup>4</sup>LAL, Université Paris-Sud and CNRS/IN2P3, Orsay, France; <sup>5</sup>Universidad Tecnica Federico Santa Maria, Chile <sup>6</sup>Università di Genova and INFN, Genova, Italy

## ABOUT THE ATLAS EVENTINDEX

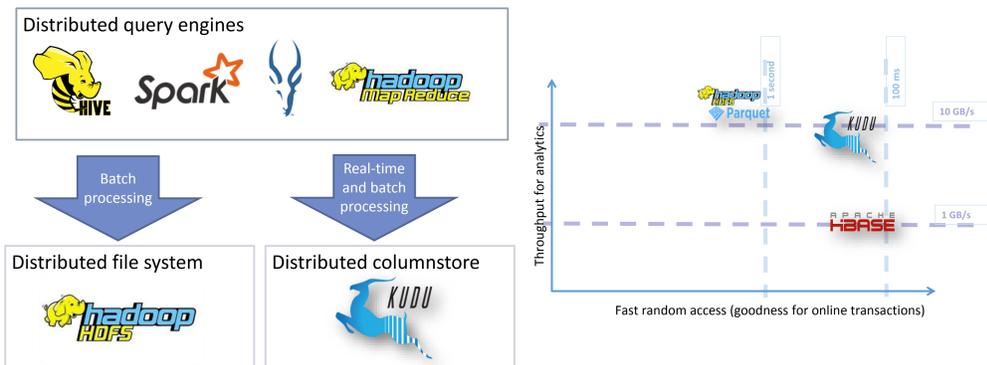
- The ATLAS EventIndex is a **catalogue** of all real and simulated **events** produced by the experiment at all processing stages [1].
- The system contains **hundreds of billions** of event records (180 billions of records as of June 2018), each consisting of ~1000 bytes.
- The goal of the ATLAS EventIndex is to allow **fast** and **efficient selection** of events of interest, based on various criteria, and provide references that point to those events in millions of files scattered in a world-wide distributed computing system.

## MOTIVATION FOR EVENTINDEX EVOLUTION

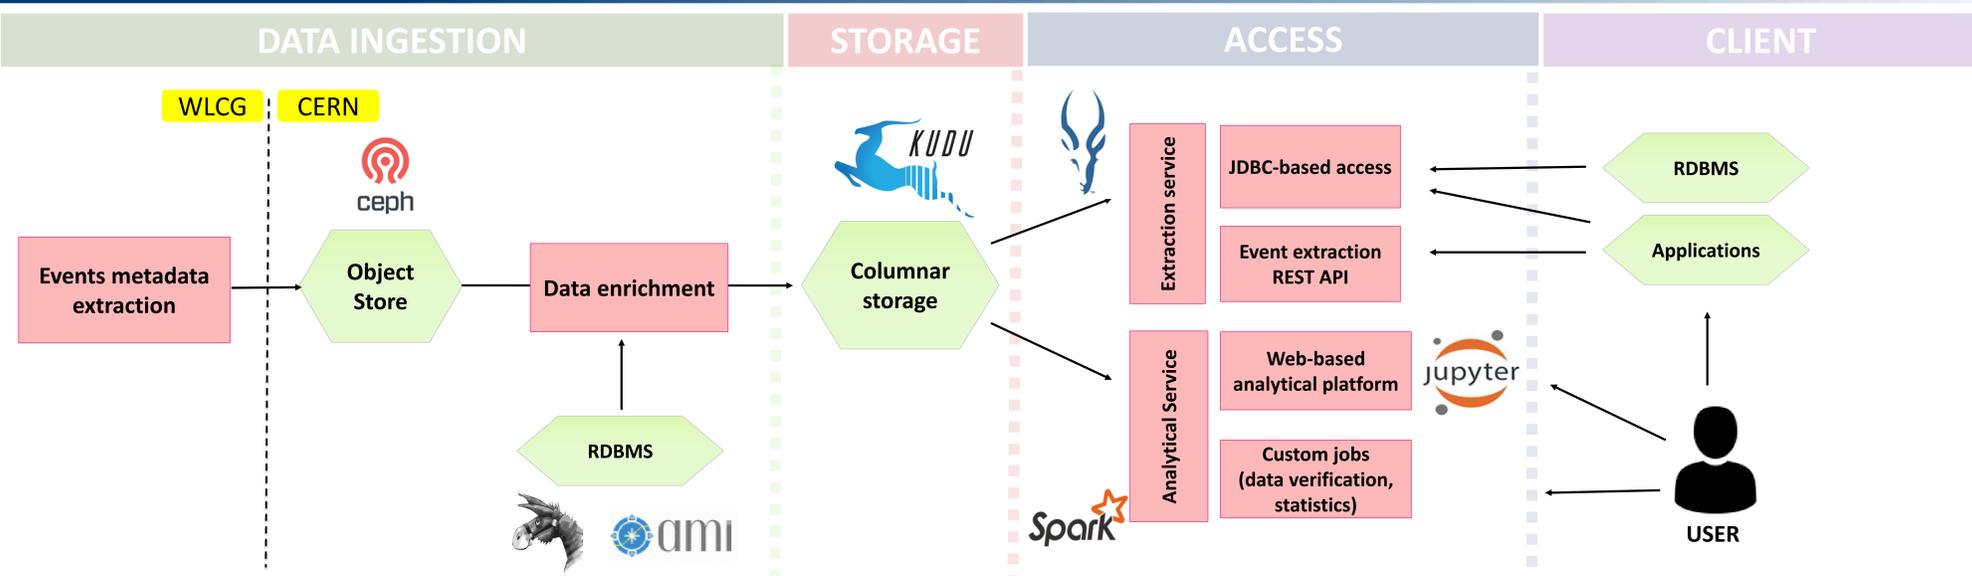
- Current EventIndex was designed in 2012-2013 using best BigData technology available at that time (Hadoop), implemented in 2014 using MapFiles and HBase, in operation since 2015 with satisfactory results [2].
- Use cases **extended** in the meantime from event picking and production completeness checks to trigger overlap studies, duplicate event detection and derivation streams (offline triggers) overlaps.
- Fast data querying based on traditional relational database involving a subset of information for real events only [3]
- Also event **rate increased** steadily throughout Run 2.
- BigData **technologies advanced** in the meantime and now we have the choice between many different products and options.
- Studies of **new data formats** and/or new storage technologies [4] concluded that Kudu is the most promising technology for the next few years. Hence this prototype.

## APACHE KUDU

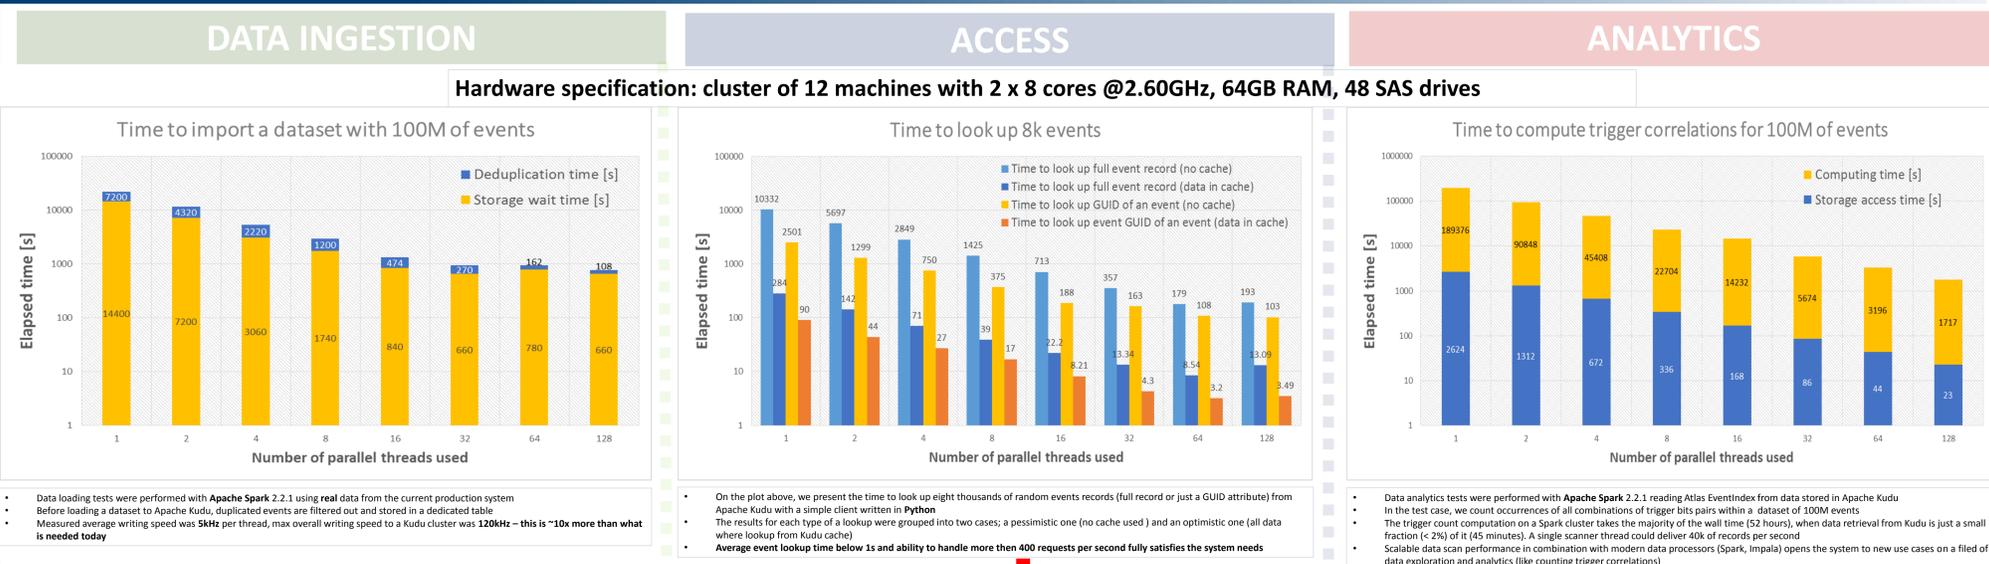
- Next generation scalable and distributed table-based storage designed for HTAP systems – **Hybrid Transactional and Analytical Processing** [5]
- Unlike Hadoop Distributed File System (HDFS), Kudu provides indexing and columnar data organization natively – this is to establish a good compromise between random **data lookups** and **analytics** performance
- Organization of the data in sharded tables with named columns, types and a primary index makes Kudu very attractive for systems with relational data models that needs to scale-out
- Apache Kudu is supported by top open-source frameworks for parallel data processing and computation including
  - Apache Spark, Apache Impala, Apache Hive, MapReduce,...



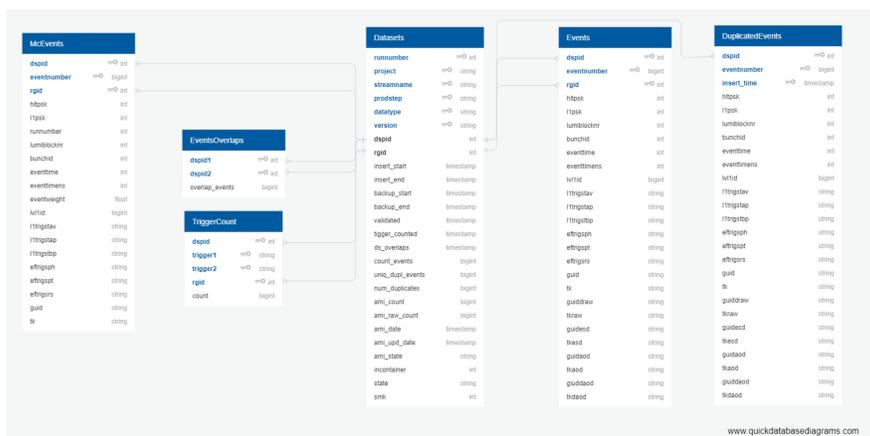
## CONCEPT OF THE NEW ATLAS EVENT INDEX PLATFORM



## MEASURED PERFORMANCE WITH APACHE KUDU STORAGE



## DATA SCHEMA PROTOTYPE IN APACHE KUDU



## REFERENCES

- Barberis D et al. 2014 The ATLAS EventIndex: an event catalogue for experiments collecting large amounts of data, J. Phys. Conf. Ser. 513 042002
- Barberis D et al. 2015 The ATLAS EventIndex: architecture, design choices, deployment and first operation experience, J. Phys. Conf. Ser. 664 042003
- Gallas E et al. 2016 an Oracle-based Event Index for ATLAS
- Z. Baranowski et al 2017 J. Phys.: Conf. Ser. 898 062020
- Lipcon T et al., 2015 Kudu: Storage for fast analytics on fast data.