Data Analytics and CERN IT Hadoop Service

CERN openlab, INTEL visit
CERN, February 2017
Luca Canali, IT-DB

Data Analytics at Scale – The Challenge

- When you cannot fit your workload in a desktop
  - Data analysis and ML algorithms over large data sets
  - Deploy on distributed systems
- Complexity quickly goes up
  - Data ingestion tools and file systems
  - Storage and processing engines
  - ML tools that work at scale

Engineering Effort for Effective ML

- From “Hidden Technical Debt in Machine Learning Systems”, D. Sculley at al. (Google), paper at NIPS 2015

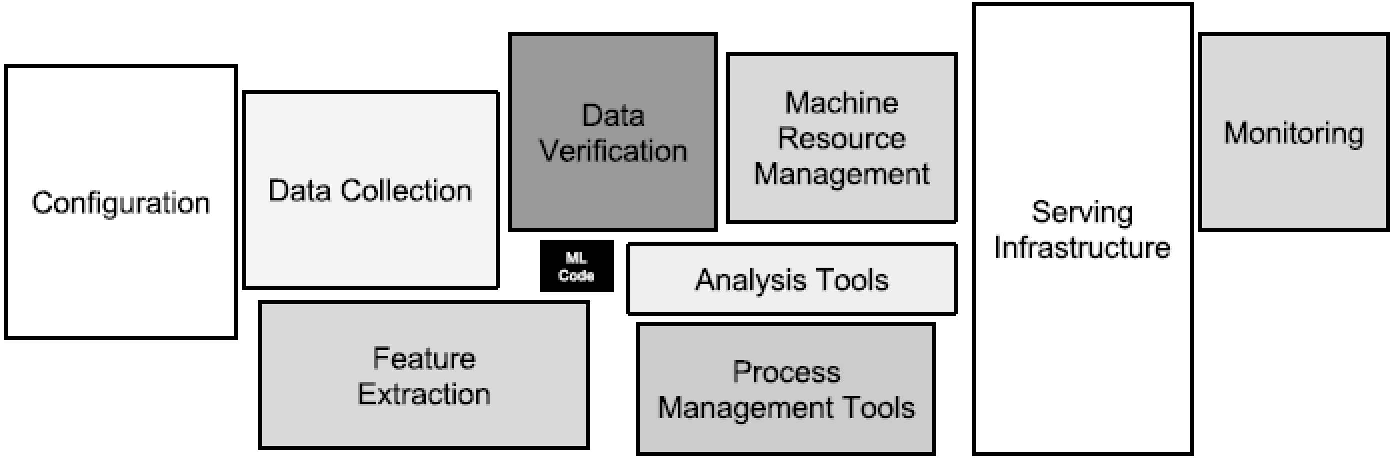

Figure 1: Only a small fraction of real-world ML systems is composed of the ML code, as shown by the small black box in the middle. The required surrounding infrastructure is vast and complex.

Managed Services for Data Engineering

- Platform
  - Capacity planning and configuration
  - Define, configure and support components
- Running central services
  - Build a team with domain expertise
  - Share experience
  - Economy of scale

Hadoop Service at CERN IT

- Setup and run the infrastructure
- Provide consultancy
- Build user community
- Joint work
  - IT-DB and IT-ST

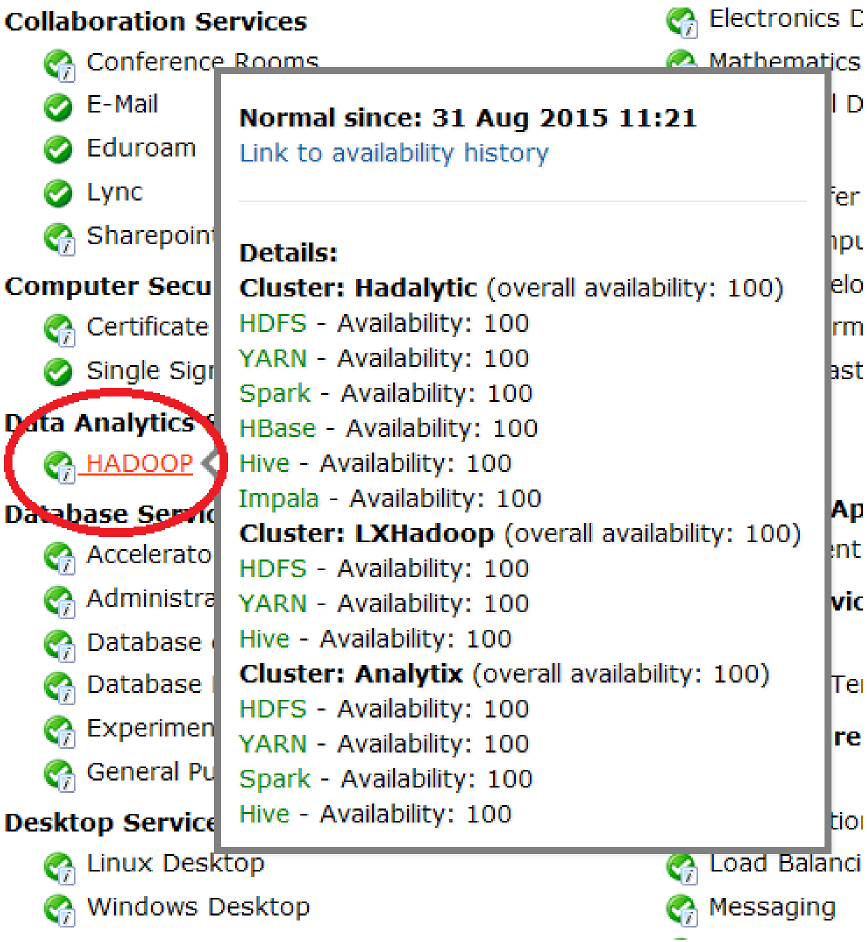

Overview of Available Components (Dec 2016)

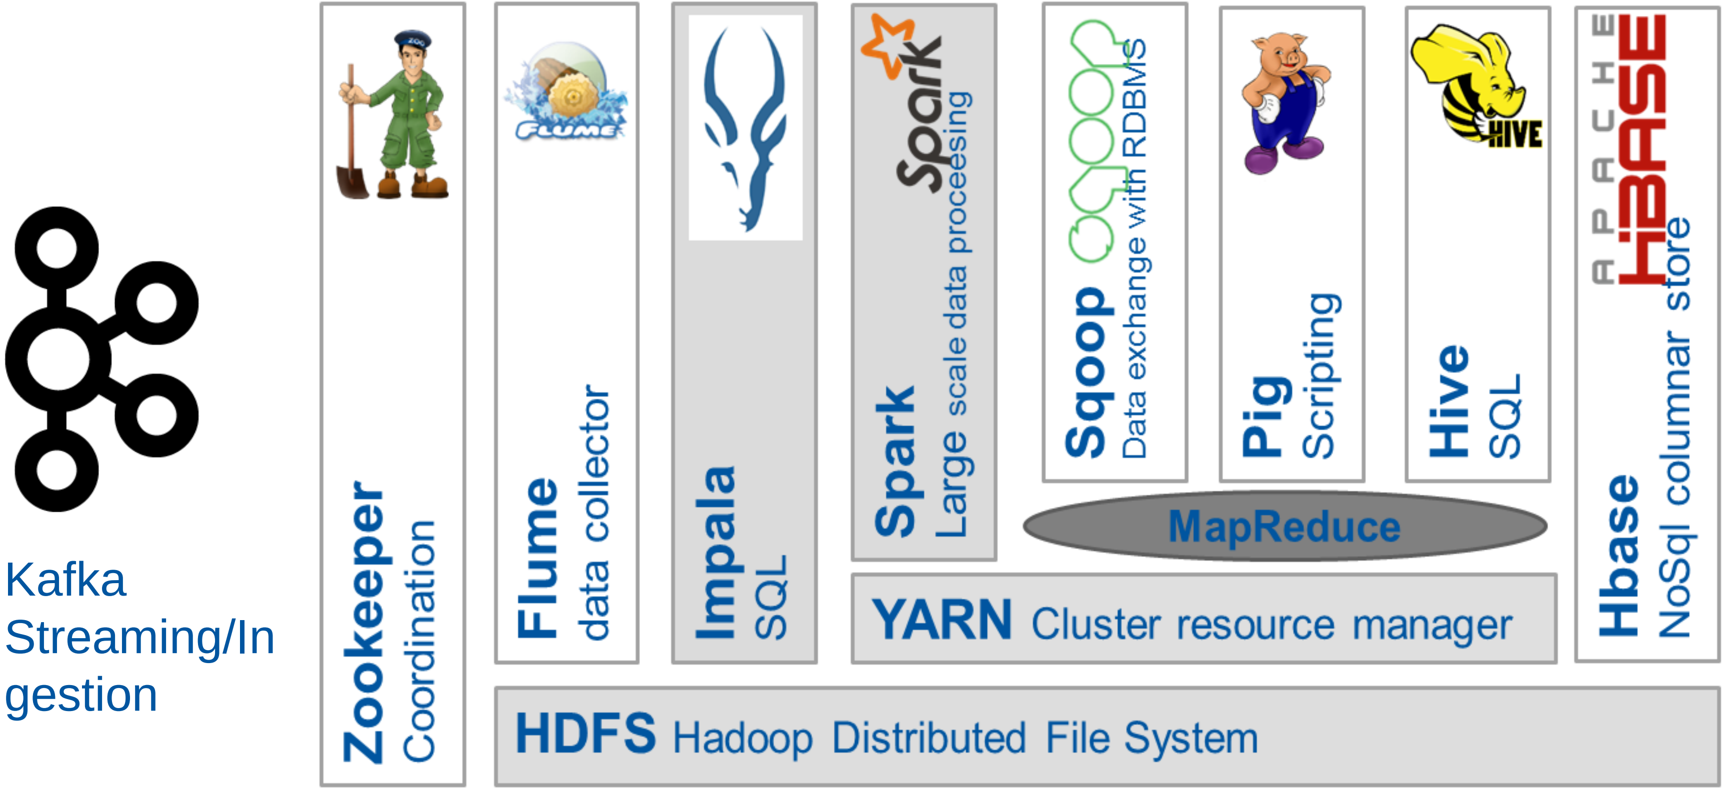

Hadoop clusters at CERN IT

- 3 production clusters (+ 1 for QA) as of December 2016
  - In the pipeline for 2017 new system for BE NXCALs platform

Cluster Name	Configuration	Primary Usage
lxhadoop	22 nodes (cores – 560,Mem – 880GB,Storage – 1.30 PB)	Experiment activities
analytix	56 nodes (cores – 780,Mem – 1.31TB,Storage – 2.22 PB)	General Purpose
hadalytic	14 nodes (cores – 224,Mem – 768GB,Storage – 2.15 PB)	SQL-oriented engines and datawarehouse workloads

Apache Spark

- Spark evolution from map reduce ideas
- Powerful engine, in particular for data science and streaming
  - Aims to be a “unified engine for big data processing”

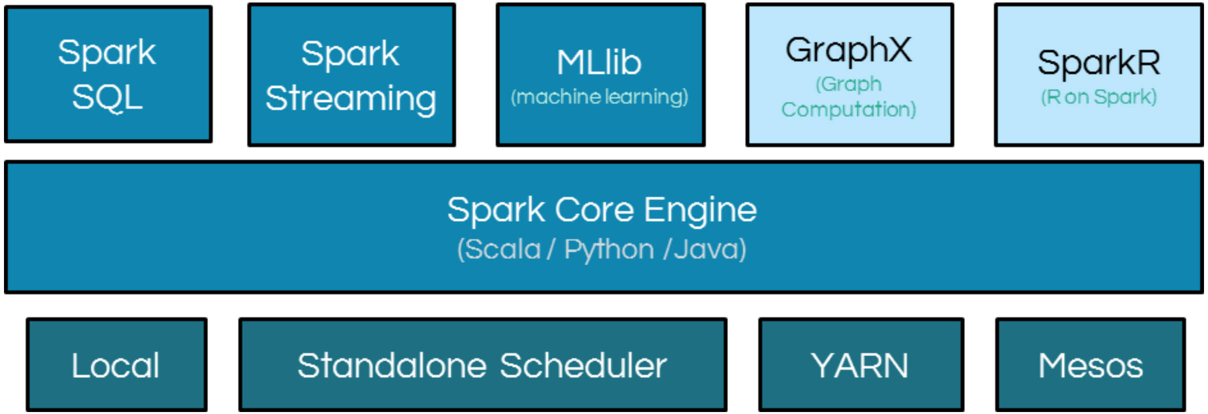

Some Important Use Cases

- Accelerator logging
- Industrial controls
- Analytics on monitoring data
- Physics analysis
  - Development of Big Data solutions for physics

Next Gen. Archiver for Accelerators Logs

Pilot architecture tested by CERN Accelerator Logging Services
Critical system for running LHC - 500 TB today, growing 200 TB/year
Challenge: production in 2017

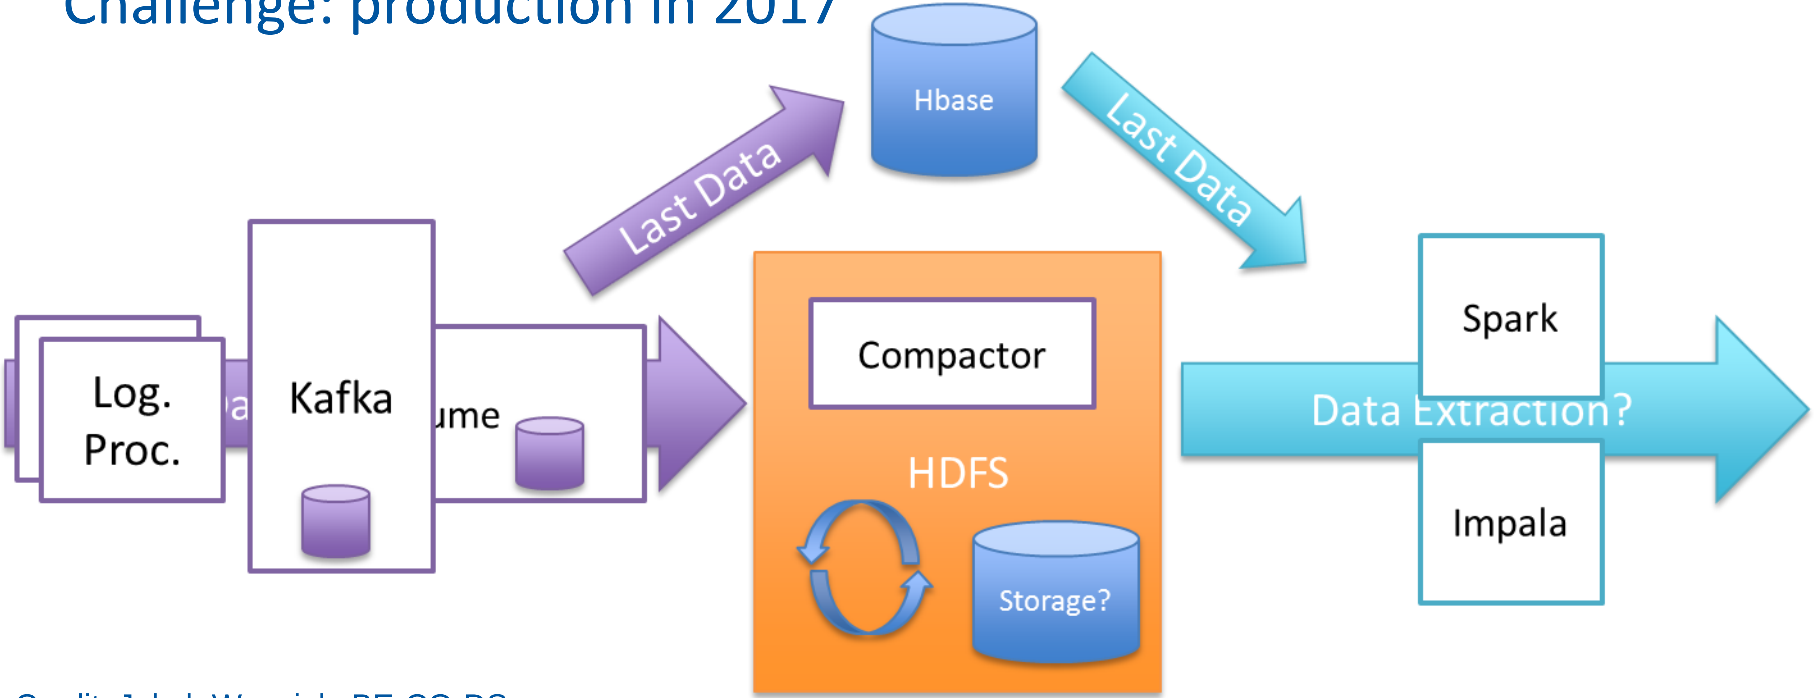

Credit: Jakub Wozniak, BE-CO-DS

Industrial Controls Systems

- Development of next generation archiver
- Currently investigating possible architectures (openlab project)
  - Including potential use of Apache Kudu

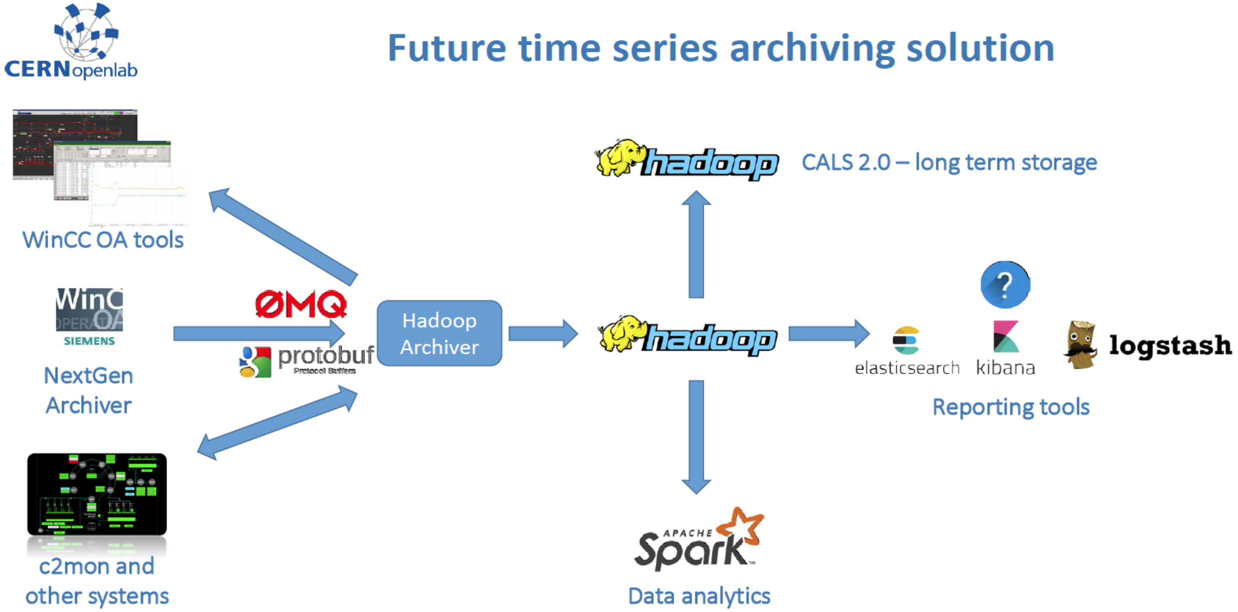

Credits: CERN BE Controls team

Analytics platform for controls and logging

- Use distributed computing platforms for storing analyzing controls and logging data
  - Scale of the problem 100s of TBs
- Build an analytics platform
  - Technology: focus on Apache Spark
  - Empower users to analyze data beyond what is possible today
  - Opens use cases for ML on controls data

Production Implementation – WLCG Monitoring

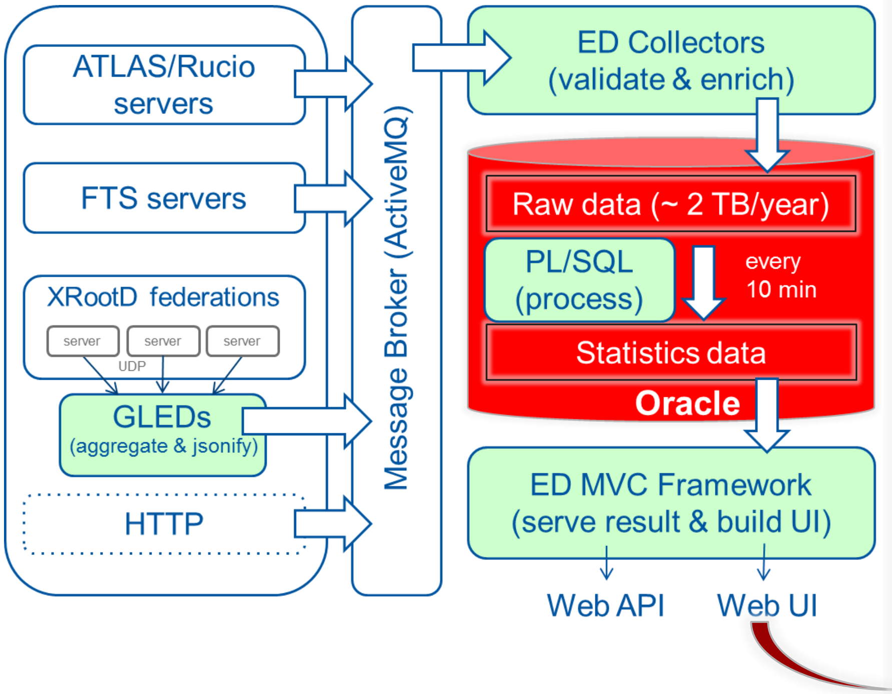

Lambda Architecture

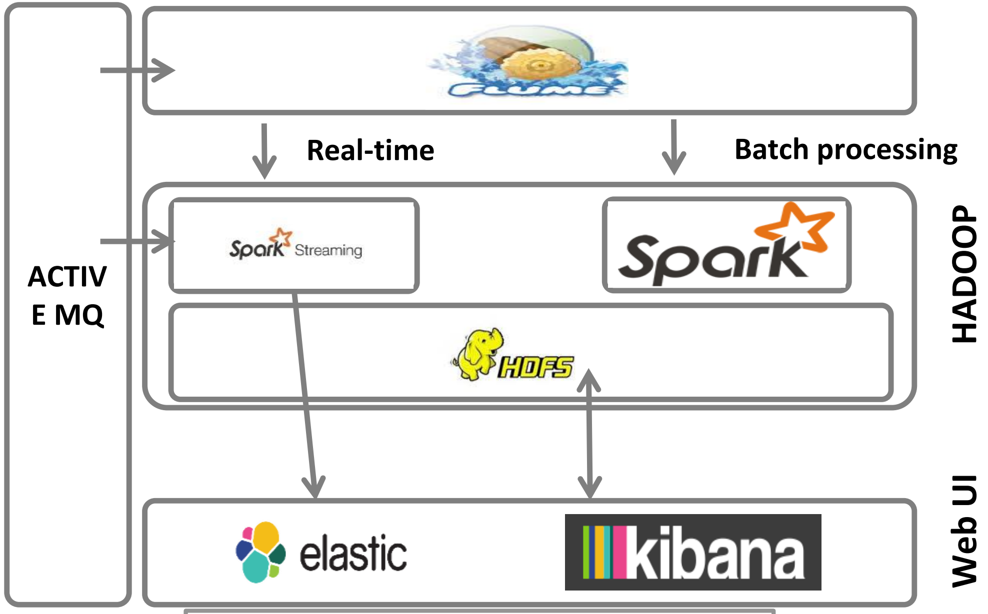

Modernizing the applications for ever demanding analytics needs

Jupyter Notebooks

- Jupyter notebooks for data analysis
  - System developed at CERN (EP-SFT) based on CERN IT cloud
  - SWAN: Service for Web-based Analysis
  - ROOT and other libraries available

- Integration with Hadoop and Spark service
  - Distributed processing for ROOT analysis
  - Access to EOS and HDFS storage

CMS Big Data Project and Openlab

Proposal: CMS Data Reduction Facility

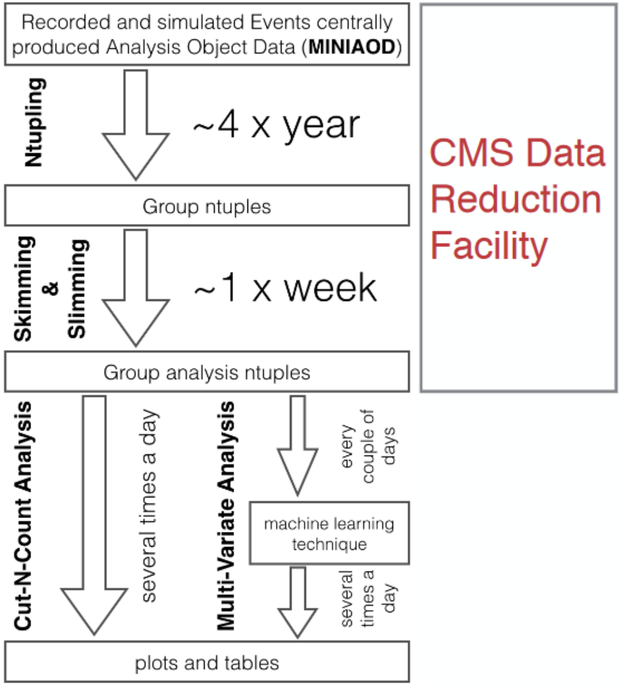

- Demonstration facility optimized to read through petabyte sized storage volumes
  - Produce sample of reduced data based on potentially complicated user queries
  - Time scale of hours and not weeks
- If successful, this type of facility could be a big shift in how effort and time is used in physics analysis
  - Same infrastructure and techniques should be applicable to many sciences

Physics Analysis and "Big Data" ecosystem

- Challenges and goals:
  - Use tools from industry and open source
    - Current status: Physics uses HEP-specific tools
    - Scale of the problem 100s of PB – towards hexascale
  - Develop interfaces and tools
    - Already developed first prototype to read ROOT files into Apache Spark
  - Challenge: testing at scale

Performance and Testing at Scale

- Challenges with ramping up the scale
  - Example from the CMS data reduction challenge: 1 PB and 1000 cores
    - Production for this use case is expected 10x of that.
    - New territory to explore
- HW for tests
  - CERN clusters + external resources, example: testing on Intel Lab equipment (16 nodes) in February 2017

Machine Learning and Spark

- Spark addresses use cases for machine learning at scale
- Distributed deep learning
  - Working on use cases with CMS and ATLAS
  - Custom development: library to integrate Keras + Spark
  - Testing also other solutions (BigDL?)
  - Room to test HW: for example FPGAs vs. GPUs etc

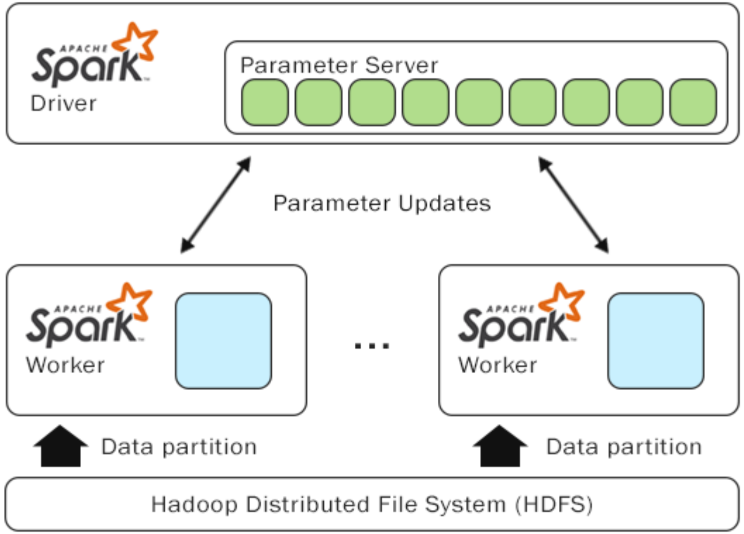

Acknowledgements

The following have contributed to the work reported in this presentation

- Members of IT-DB and IT-ST
  - Supporting Hadoop core and components