Data Analytics – Use Cases, Platforms, Services

ITMM, March 5th, 2018

Luca Canali, IT-DB

Analytics and Big Data Pipelines – Use Cases

- Many use cases at CERN for analytics
  - Data analysis, dashboards, plots, joining and aggregating multiple data, libraries for specialized processing, machine learning, …
- Communities
  - Physics:
    - Analytics on computing data (e.g. studies of popularity, grid jobs, etc) (CMS, ATLAS)
    - Development of new ways to process ROOT data, e.g.: data reduction and analysis with spark-ROOT by CMS Bigdata project, also TOTEM working on this
  - IT:
    - Analytics on IT monitoring data
    - Computer security
  - BE:
    - NX CALS – next generation accelerator logging platform
    - BE controls data and analytics
  - More:
    - Many tools provided in our platforms are popular and readily available, likely to attract new projects
    - E.g. Starting investigations on data pipelines for IoT (Internet of Things)

“Big Data”: Not Only Analytics

- Data analytics is a key use case for the platforms
- Scalable workloads and parallel computing
  - Example work on data reduction (CMS Big Data project) and parallel processing of ROOT data (EP-SFT)
- Database-type workload also important
  - Use Big Data tools instead of RDBMS
  - Examples: NXCALS, ATLAS EventIndex, explorations on WINCC/PVSS next generation
- Data pipelines and streaming
  - See example of monitoring and Computer security (Kafka development with help of CM)
  - Also current investigations on IoT (project with CS)

New IT Monitoring

Critical for CC operations and WLCG

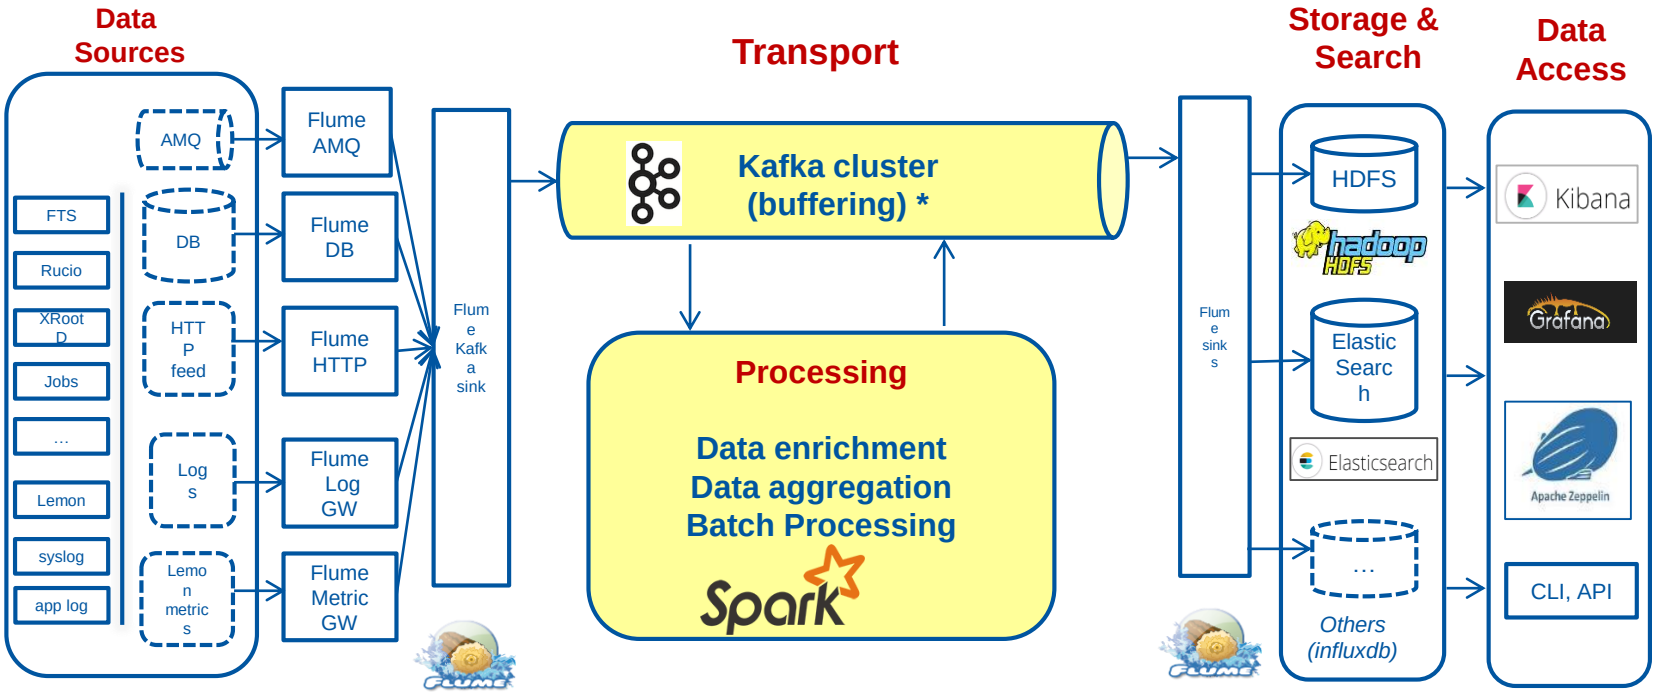

- Data now 200 GB/day, 200M events/day
- At scale 500 GB/day
- Proved effective in several occasions

Credits: Alberto Aimar, IT-CM-MM

Computer Security intrusion detection use cases

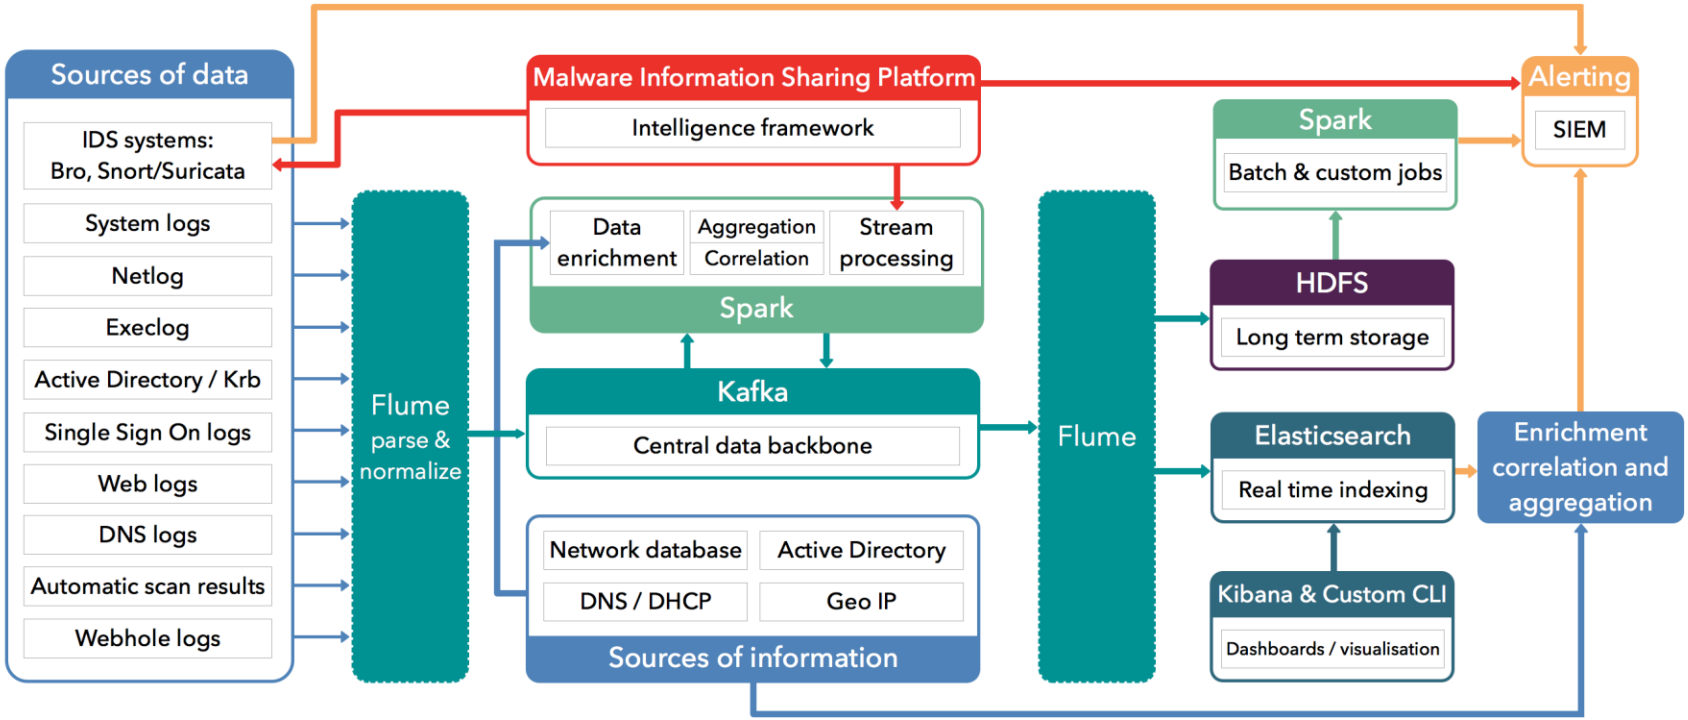

Credits: CERN security team, IT-DI

Example, BE-NXCALS

Many parts

- Streaming
- Online system
- APIs for data extraction and analytics

Note:

This will replace and enhance current platform based on RDBMS

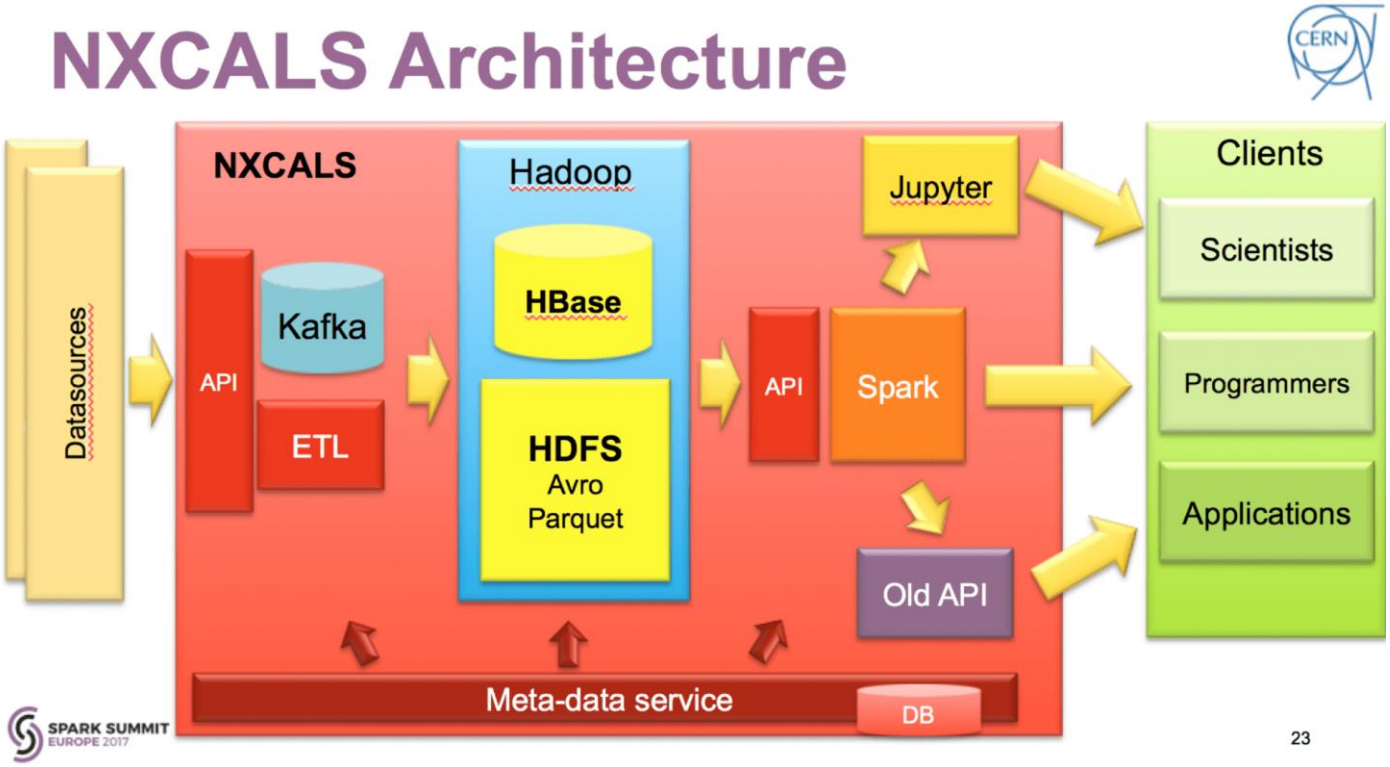

ATLAS EventIndex

- Searchable catalog of ATLAS events
  - First “Big Data” project in our systems
  - Over 80 billions of records, 140TB of data

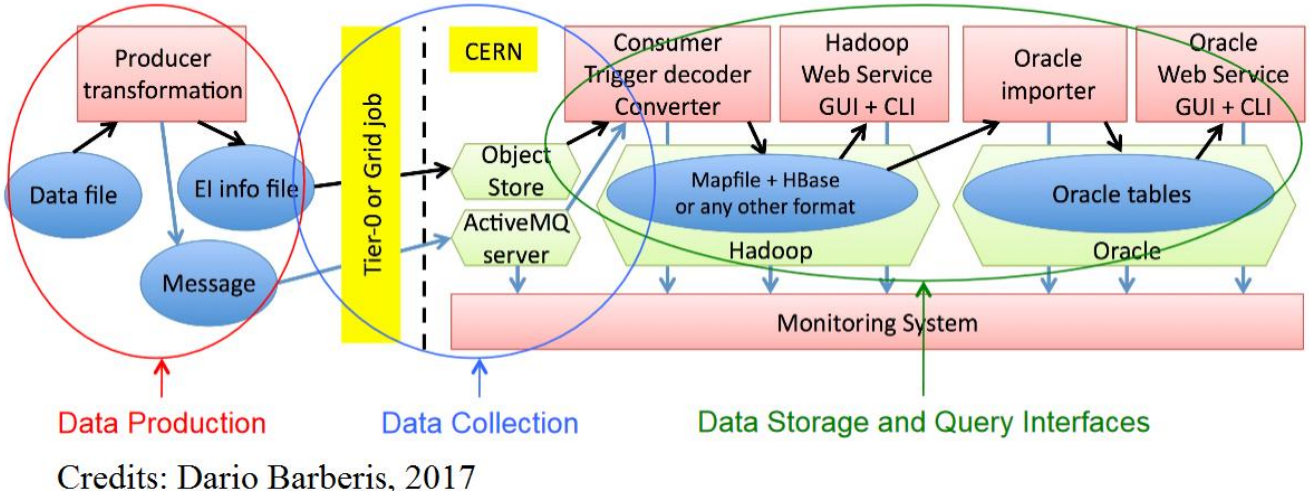

Credits: Dario Barberis, 2017

Goal: Meet Use Cases With Managed Services for Data Engineering and Analytics

- Platform
  - Capacity planning and configuration
  - Define, configure and support components
  - Use tools and ideas from “Big Data”/ open source communities
- Running central services
  - Build a team with domain expertise
  - Share experience, provide consultancy
  - Economy of scale

Highlights of “Big Data” Components

- Apache Hadoop clusters with YARN and HDFS
  - Also HBase, Impala, Hive,…
- Apache Spark for analytics
  - Apache Kafka for streaming
- Data: Parquet, JSON, ROOT
- UI: Notebooks/ SWAN

Tools and Platforms – Ecosystem is Large

- Many tools available: from HEP + other sciences + open source
  - Challenges and opportunities: build your solution + expect natural evolution
  - Components are modular, users adoption and traction drives efforts
- Front ends

- Engines and Systems

- Storage and data sources

Databases

- Data formats: ROOT, Parquet, JSON,..

Hadoop Clusters at CERN IT

- 3 current production clusters
  - + environments for NXCALS DEV and HadoopQA
- A new system for BE NXCALs (accelerator logging) platform
  - Currently being commissioned

Cluster Name	Specification	Primary Usage
lxhadoop	18 nodes (cores - 520,Mem - 701GB,Storage - 1.29 PB)	ATLAS EventIndex
analytix	40 nodes (cores - 1224,Mem - 4.11TB,Storage - 4.32 PB)	General Purpose
hadalytic	12 nodes (cores - 364,Mem - 580GB,Storage - 2.081 PB)	Development Cluster (old BE prod)
NxCALS	24 nodes (cores - 1152,Mem - 12TB,Storage - 4.6PB,SSD - 92 TB)	Accelerator Logging Service

Workload Metrics

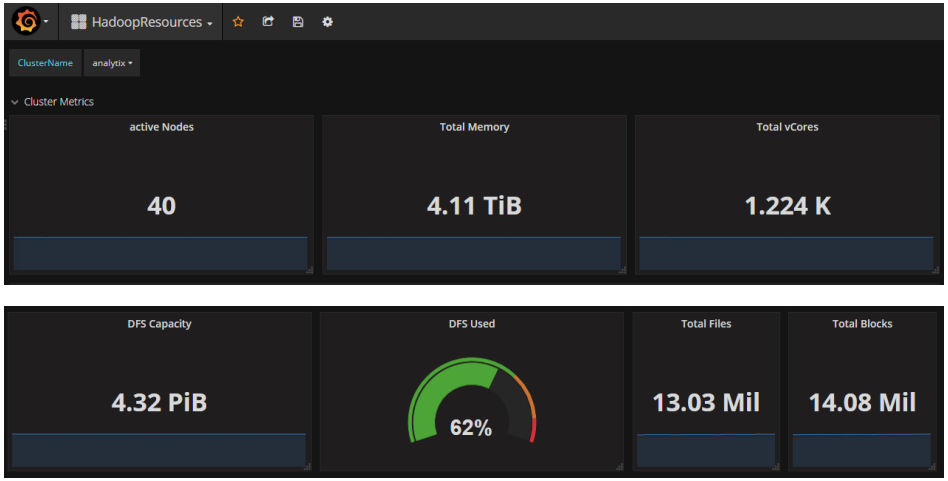

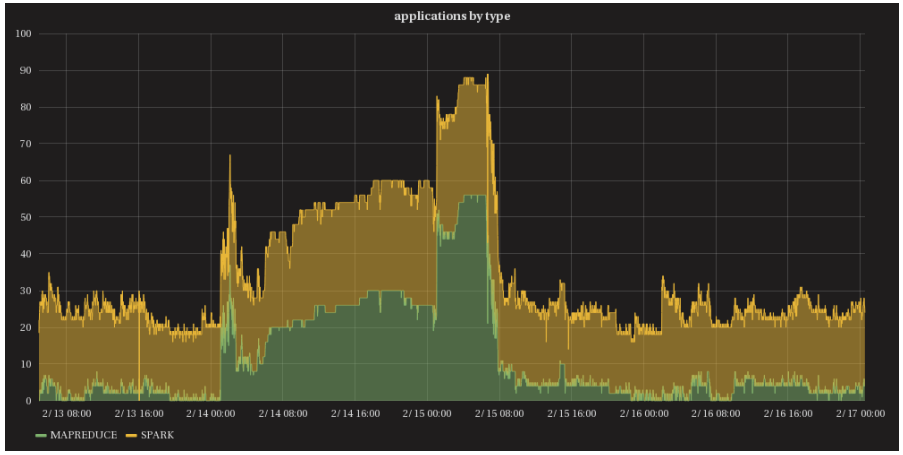

Projects and Activities

- Spark integration with SWAN (Jupyter notebook service)
- Spark for HEP (CMS Big Data, Spark and EOS)
- Kafka Service and Self-service tools
- Kudu: online and analytics, best of HBase and HDFS?
- Spark as a service (Spark and Kubernetes)
- Courses, consultancy, forums, workshops and conferences, documentation, etc

Spark for Physics

- Promising results
  - See ACAT 2017: “CMS Analysis and Data Reduction with Apache Spark”
  - Processing ROOT data using Big Data tools (Spark)
  - Aligned with the goals of the HEP software foundation activities
- Currently working with CMS Bigdata and Intel on openlab
  - Developed Hadoop-XRootD connector (read EOS from Spark)
  - Working on large-scale reduction and analysis
- Also collaborating with Totem

Spark as a Service and Cloud

- Exploring Spark on Kubernetes
  - Separate compute cluster from storage
  - Already used in industry, now mainstream with latest Spark
  - Exploration phase for us at this stage
- Motivations
  - Can this help to scale out on large workloads and help streamline resource allocations?
- Context
  - For the use case of using Spark to process physics data on EOS, running Spark on Hadoop with HDFS has limited advantages
  - Spark is basically a library and needs a cluster manager: YARN/Hadoop, Mesos, Stand-alone, since Q1 2018 also Kubernetes

Spark Integration with SWAN – Towards an Analytics Platform

- Goal: become go-to platform for analytics at CERN
  - Lower the barrier to start doing analytics and provide a pre-configured environment with the most used tools
- Integration of SWAN with Spark clusters in IT
  - Key component for BE-NXCALS: chosen as user interface
  - Already used by several pilot users too, for example for monitoring data analytics
  - opening to all users later in 2018. Currently SWAN default for Spark is to use “local mode” (no cluster)
  - Joint efforts with EP-SFT, IT-ST, IT-CM on integrating functionality, security and monitoring aspects

Jupyter Notebooks and Analytics Platforms

Apache Spark

LHCb Opendata

This notebok illustrates a nice interplay between LHCb Opendata, Spark and ROOT - all teaching interesting flavour Physics!

Plotting a feature:

You can plot any feature of the data in a histogram. Choose any suitable binning that allows you to observed the distribution of the variable clearly. I...ing a histogram for the first kaon candidate's momentum x-component (H1_PX):

Open in SWAN

In [8]:
```
# Plot a histogram of the distribution of the H1_PX variable, using Pandas

# This is a basic solution that  moves all the data from the Spark DataFrame
# into a Python Pandas DataFrame. It's OK for small sata sets, but it has scalability issues

h1px_data = sim_data_df.select("H1_PX").toPandas() # select H1_PX data and moves it to Pandas
h1px_data.plot.hist(bins=31, range=[-150000, 150000], title="Histogram - distribution of H1_PX, simulation data")
xlabel('H1_PX (MeV/c)')
ylabel('Count');
```

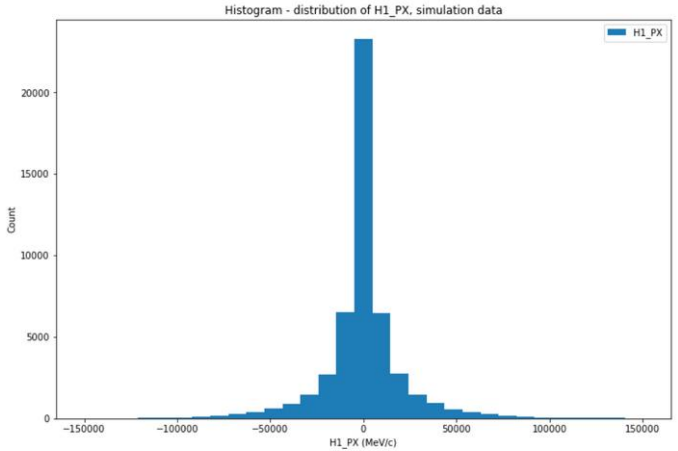

SWAN Customisation

Specify the parameters that will be used to contextualise the container which is created for you. See the online SWAN guide for more details.

Software stack more...

89

Platform more...

x86_64-slc6-gcc49-opt

Environment script more...

e.g. $CERNBOX_HOME/MySWAN/myscript.sh

Number of cores more...

Number of cores to associate to the container.

1

Spark cluster more...

Hadalytic

Start my Session

Jupyter Notebooks and Analytics Platforms

Bringing Database Workloads to Big Data

- Goal: migrations from RDBM to “Big Data”
  - Technically challenging: need systems that can perform “Hybrid transactional/analytical processing” (HTAP)
  - Not all workloads can fit, but when applicable can provide gain in performance and reduced cost
- Technology exploration:
  - Apache Kudu: a storage system that can be used with Big Data tools, with Indexes and columnar features
- Testing: next gen ATLAS Event Index, next gen WINCC (PVSS) archiver

Users and community

- Develop value of the services for users and community
  - We have developed and delivered short courses on Spark and Hadoop
    - More session(s) scheduled in 2018 + planning course on Kafka
  - Plannig -> presentations at relevant forums (ASDF, IT Tech, ITUM)
  - Consultancy and projects with relevant communities
- Linking with communities outside CERN
  - Contacts with industry and open source communities
    - (Intel openlab, conferences, etc)
  - Others in HEP: Already contacted selected T1s and other centers (Sara, BNL, Princeton) in 2017 for sharing ideas on Hadoop config + CMS Bigdata (Fermi Lab, Princeton, Padova) + TOTEM (team in Warsaw)
  - Potential: contacts with ”other sciences” (astrophysics, biology/medical)

Challenges

- Platforms
  - Provide evolution for HW – currently plan to do “rolling upgrades” of ~ 25% of the platform yearly, some of the current HW is quite old
- Service
  - Build robust service for critical platform (NXCALS and more) using custom-integrated open source software solutions in constant evolution
  - Support production services (IT monitoring, Security, EventIndex)
  - Evolve service configuration and procedures to fulfil users needs
  - Further grow value for community and projects
- Knowledge and experience
  - Technology keeps evolving, need to learn and adapt quickly to change

Conclusions

- Hadoop, Spark, Kafka services
  - Unified support, engineering and consulting
  - Started work for specific use cases, experience led to opening central services to more communities (Hadoop and Spark, Spark on SWAN)
- Production service support
  - BE-NXCALS, critical service starting in 2018
  - Support for IT-Monitoring, Security and ATLAS EventINdex
  - Still developing and streamlining services (configuration, monitoring, backup, tools, etc)
- Growing value for user communities
  - Developing projects
    - Openlab with CMS Big Data and Intel for Physics analysis and data reduction with Spark
    - Evolution of WINCC/PVSS and Atlas EventIndex
  - Leverage experience and bring solutions around Big Data to a growing community
    - Lower barrier to entry (Spark on notebook), organize courses, presentations, and knowledge sharing activities in general