

Big Data, Information Discovery and the biggest machine in the world

Antonio Romero Marín

Manuel Martin Marquez

Big Data Breakfast EMEA @ Oracle OpenWorld 2016



What's CERN

- CERN - European Laboratory for Particle Physics
- Founded in 1954 by 12 Countries for fundamental physics research in a post-war Europe
 - Science for Peace
- Worldwide Collaboration



Discoveries

Seeking answers to questions about the Universe



International Cooperation

Uniting bringing nations together through science



Science & Education


Training tomorrow's scientists and engineers



Science & Technology

Advancing the frontiers of technology

Fundamental Research



The Large Hadron Collider (LHC)



- Largest machine in the world
 - 27km, 6000+ superconducting magnets
- Fastest racetrack on Earth
 - Protons circulate 11245 times/s (99.9999991% the speed of light)
- 600 million collisions per second
 - Generating approximately one petabyte of data per second

The Large Hadron Collider (LHC)



- **Emptiest place** in the solar system
 - High vacuum inside the magnets
- One of the **coldest places** on Earth
 - Main magnets operate at a temperature of 1.9 K (-271.3°C)
- **Hottest spot** in the galaxy
 - During Lead ion collisions create temperatures 100000x hotter than the heart of the sun

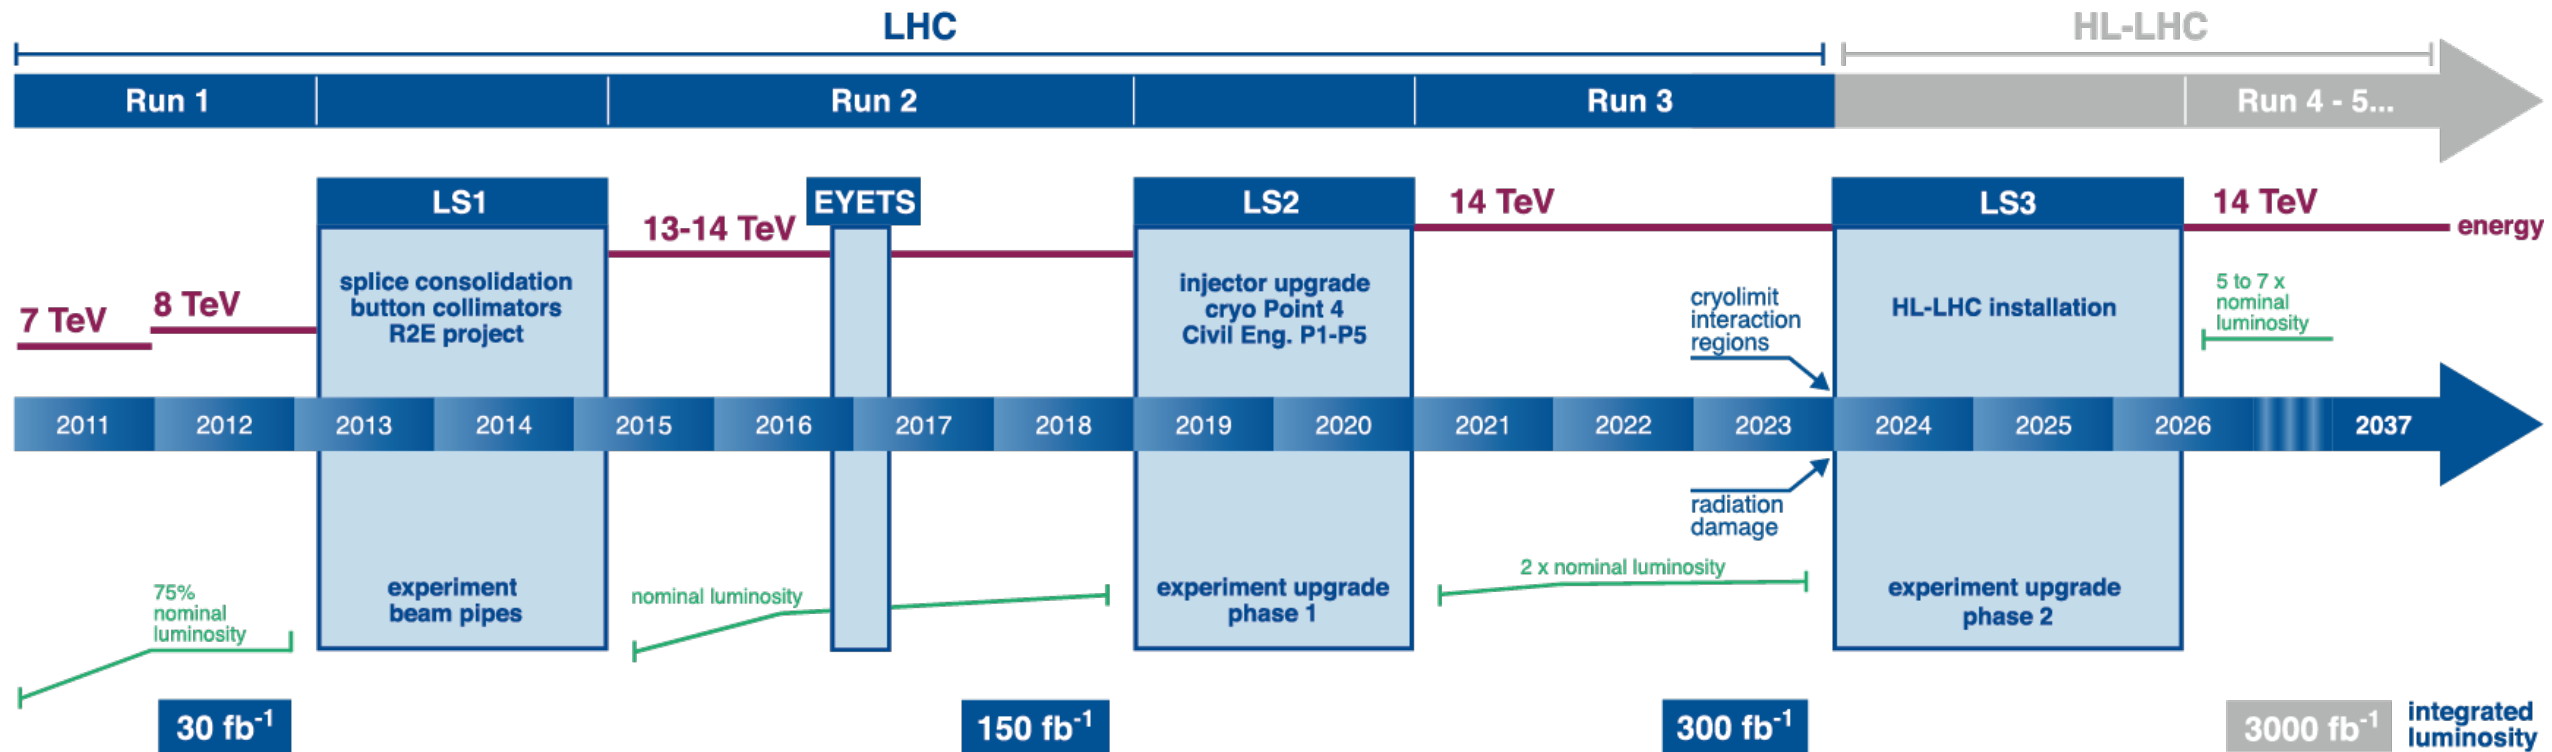
ATLAS Detector

The image shows the ATLAS detector, a large particle detector used in high-energy physics experiments. It is a complex structure with many layers of sensors and detectors. The main structure is blue, and there are several large cylindrical components with orange and white stripes. The detector is surrounded by a network of pipes, cables, and structural elements. The background shows a large industrial facility with various equipment and lighting.

150 Million of sensor
Control and detection sensors

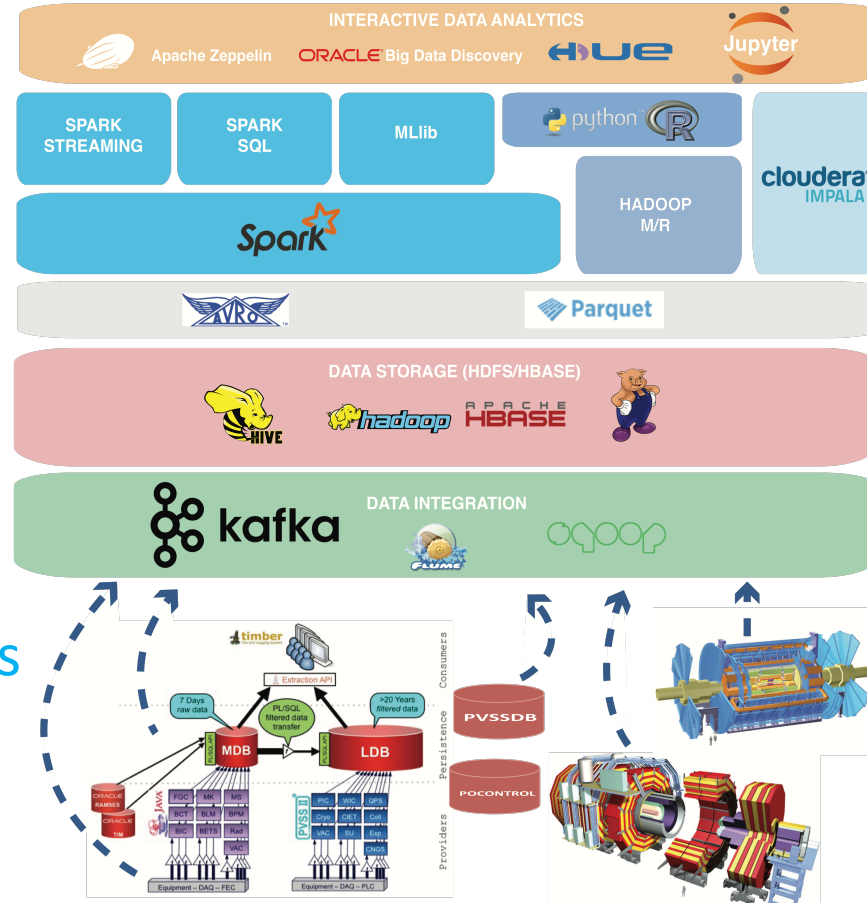
Massive 3D camera
Capturing 40+ million collisions per second
Data rate TB per second

LHC / HL-LHC Plan



Hadoop and Analytics – IT-DB-SAS

- New scalable data services
 - Scalable databases
 - Hadoop ecosystem
 - Time Series databases
- Big Data Analytics
- Activities and objectives
 - Support of Hadoop Components
 - Further value of Analytics solutions
 - Define scalable platform evolution
- Hadoop Production Service



Collaboration Services

- Conference Rooms
- E-Mail
- Eduroam
- Lync
- Sharepoint

Computer Security

- Certificate
- Single Sign-On

Data Analytics

- HADOOP**

Database Services

- Accelerator
- Administration
- Database
- Database Migration
- Experimentation
- General Purpose

Desktop Services

- Linux Desktop
- Windows Desktop

Details:

Cluster: Hadalytic (overall availability: 100)
 HDFS - Availability: 100
 YARN - Availability: 100
 Spark - Availability: 100
 HBase - Availability: 100
 Hive - Availability: 100
 Impala - Availability: 100

Cluster: LXHadoop (overall availability: 100)
 HDFS - Availability: 100
 YARN - Availability: 100
 Hive - Availability: 100

Cluster: Analyticx (overall availability: 100)
 HDFS - Availability: 100
 YARN - Availability: 100
 Spark - Availability: 100
 Hive - Availability: 100

CERN Accelerator Logging Service



- +800 extraction clients
- +5 million extraction requests per day
- 130 custom applications

~ 250'000 Signals
 ~ 50 data loading processes
 ~ 5.5 billion records per day
 ~ 275 GB / day
 → 100 TB / year throughput



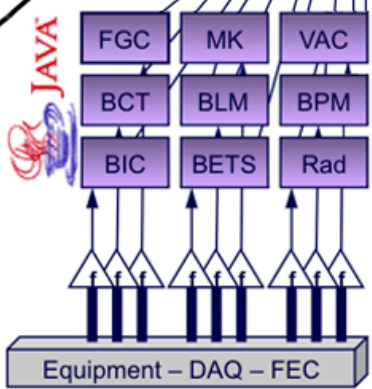
~ 1 million signals
 ~ 300 data loading processes
 ~ 4 billion records per day
 ~ 160 GB / day
 → 52 TB / year stored

7 Days raw data

PL/SQL filtered data transfer

>20 Years filtered data

Filters for data Reduction



Consumers

Persistence

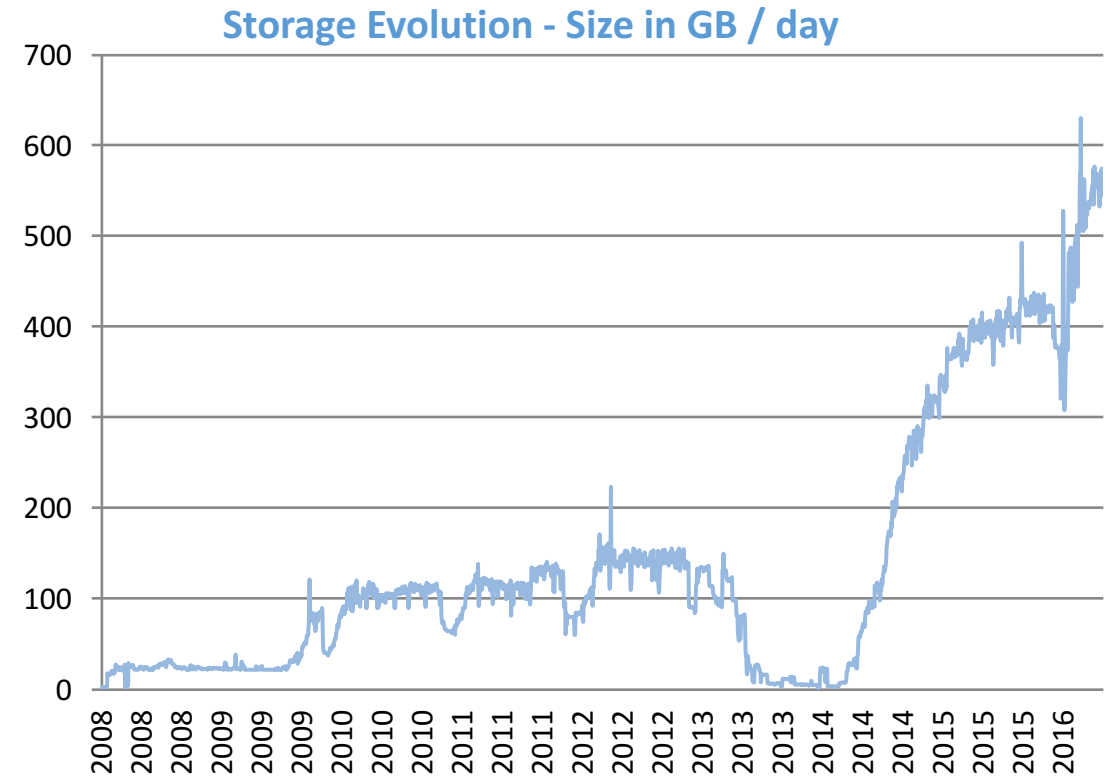
Providers

Credit: BE-CO-DS



CERN Accelerator Logging Service

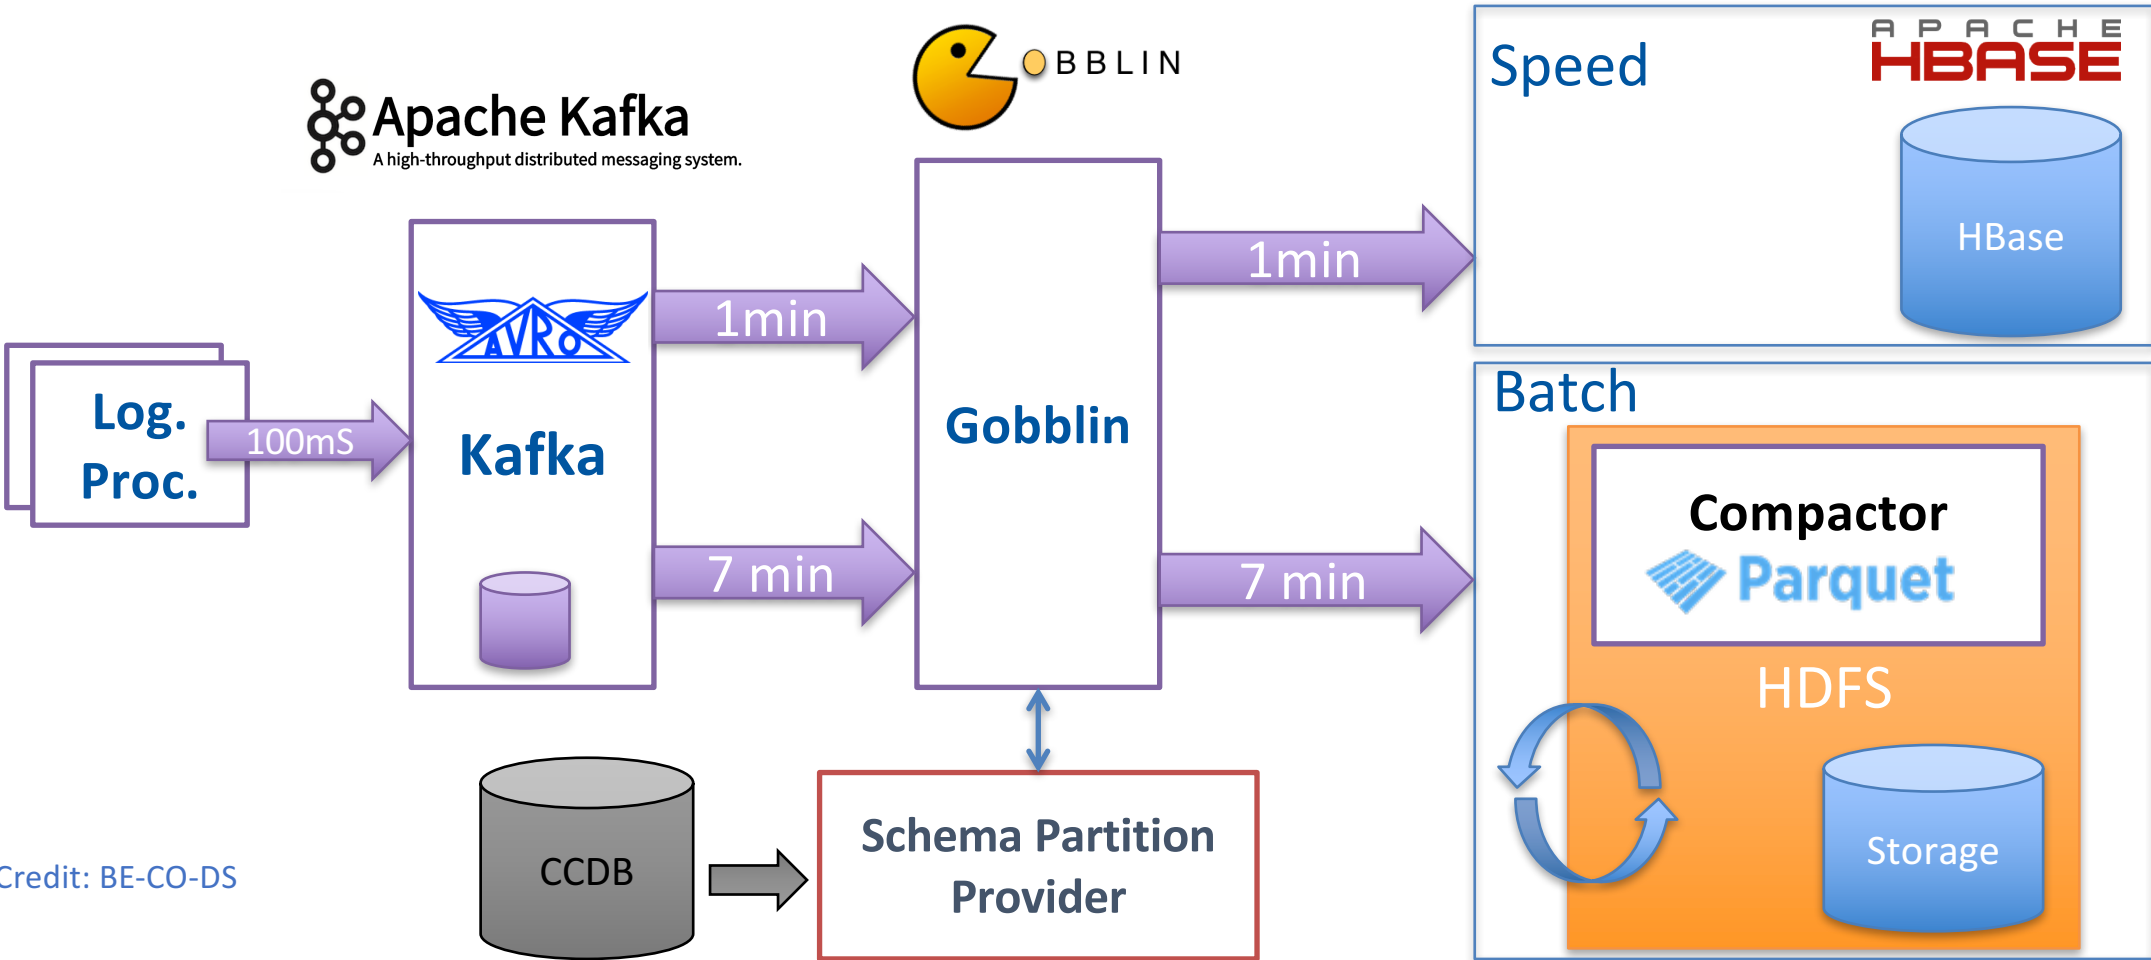
- New Landscape bring new challenges
 - Better Performance on bigger datasets
 - Big Data queries: Impala, Spark SQL
 - Leverage analytics capabilities
 - Spark Analytics: Python, ML, R
 - More heterogeneous data access models



Credit: BE-CO-DS

CERN Accelerator Logging Service

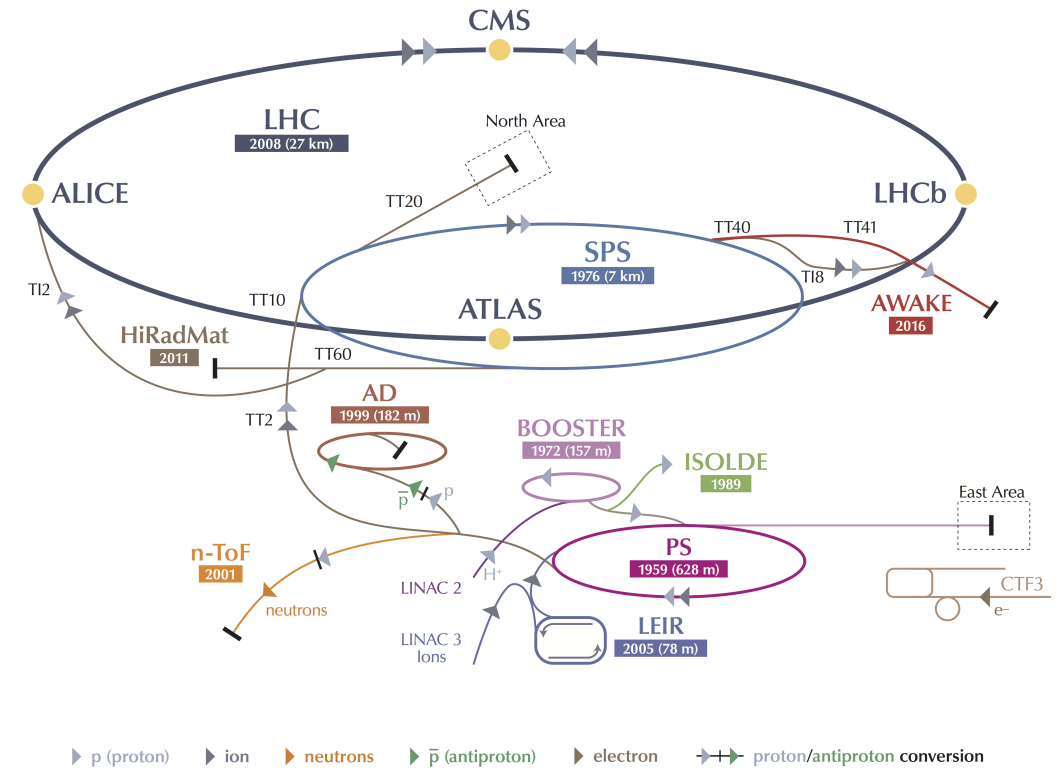
 **Apache Kafka**
A high-throughput distributed messaging system.



Credit: BE-CO-DS

Accelerator Postmortem Analysis

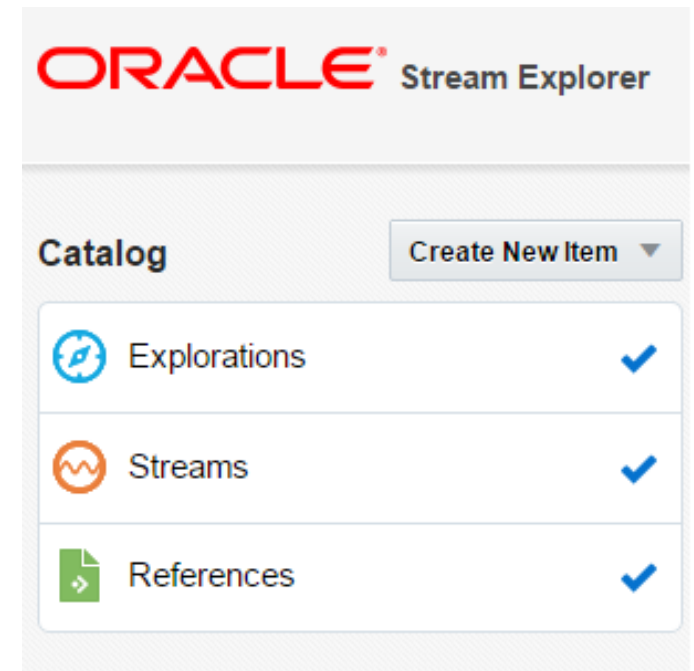
- Postmortem Analysis
 - Diagnostic on failures
 - Continue operations safely
 - Intervention Required
- Designed for CERN LHC
 - Extended to injectors complex (SPS)
 - External Post Operational Checks
 - Injection Quality Checks

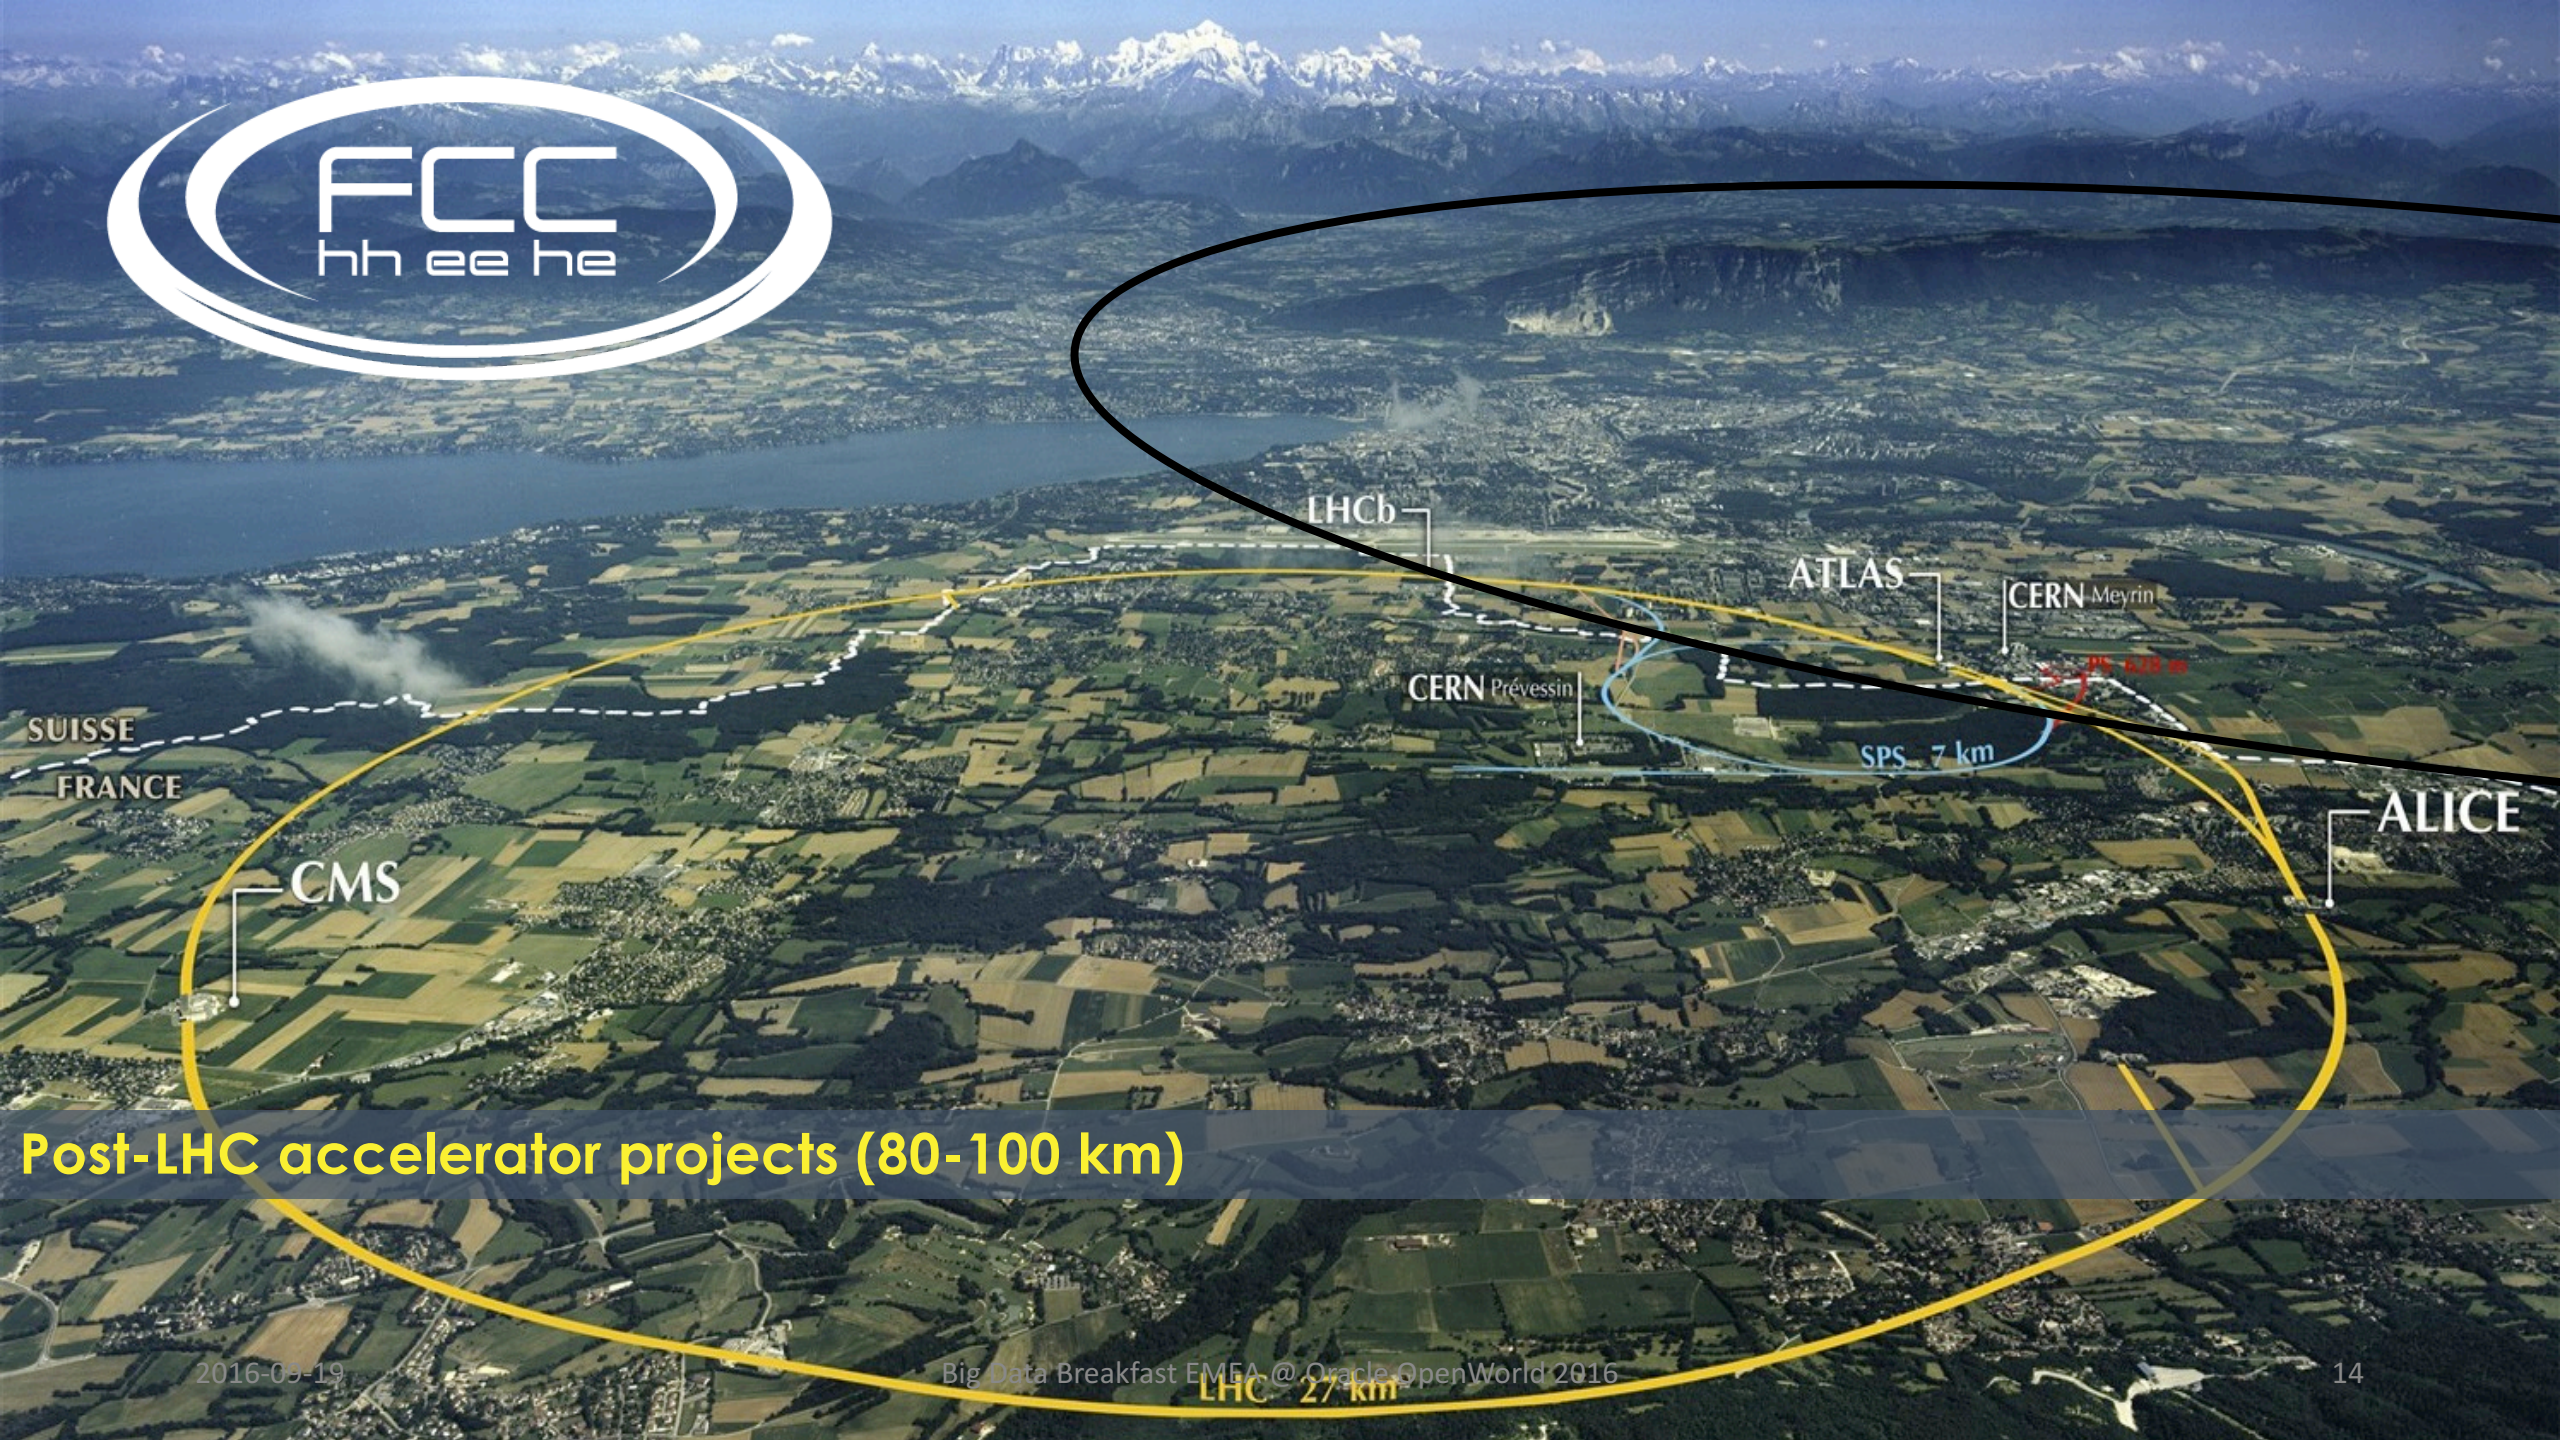


LHC Large Hadron Collider SPS Super Proton Synchrotron PS Proton Synchrotron
 AD Antiproton Decelerator CTF3 Clic Test Facility AWAKE Advanced WAKEfield Experiment ISOLDE Isotope Separator OnLine DEvice
 LEIR Low Energy Ion Ring LINAC LINear ACcelerator n-ToF Neutrons Time Of Flight HiRadMat High-Radiation to Materials

Accelerator Postmortem Analysis

- Challenges:
 - Stringent Timing Constraint
 - Better scalability
 - data storage
 - IO throughput
 - Real Big Data Streaming Analytics





LHCb

ATLAS

CERN Meyrin

CERN Prévessin

SPS 7 km

PS 6.28 km

SUISSE
FRANCE

CMS

ALICE

LHC 27 km

Post-LHC accelerator projects (80-100 km)


Use Case: FCC RAMS

- Reliability, Availability, Maintainability and Safety (RAMS) studies for the Future Circular Collider (FCC)
- Study and increase the reliability and availability of the LHC
- Use RAMS findings to assess the feasibility of the needs of FCC
- Data distributed across multiple sources
 - Operations e-logbook
 - Accelerator Fault Tracking project
 - Accelerator logging service
 - Accelerator schedules
 - Cryogenics
 - Vacuum, Power Converters, etc.




Use Case: FCC RAMS

Scenario



- LogBook
- MEDB
- CALS
- AFT
- AME

- Power Converters
- Cryogenics
- Machine Protection
- Accelerator Major Events
- Accelerator Fault Tracking
- Accelerator Logging
- Operations logbook



Deployment overview

cloudera®

CDH 5.7.1
16 nodes, 24 GB ram
Intel Xeon L5520 @ 2.27GHz
165 TB HDFS

Oracle Big Data Discovery
Libraries + Hive table detector



Resource Management (YARN)

Data Storage





Data Integration



C
o
o
r
d
i
n
a
t
i
o
n

Big Data Discovery v1.3.0
Dgraph & Studio



ORACLE®

EXALYTICS



4x Xeon E7-8895 v2 (15 cores each)
2 TB RAM
4.8 TB Flash + 6 x 1.2 TB 10K HDD

Oracle Big Data Discovery Overview

- Data Exploration & Discovery
 - Interactive catalog of all data
 - Assess attribute statistics, data quality and outliers
 - Quick data exploration or create dashboards and applications
- Data Transformation with Spark in Hadoop
 - Apply built-in transformations or write your own scripts
 - Data Enrichment
 - Text: Entity extraction, relevant terms, sentiment, language detection
 - Geographical information: address, IP, reverse
 - Preview results, undo, commit and replay transforms
- Collaborative environment
 - Share and bookmarks
 - Create and share transformed datasets

Discovery Applications

Explore ▾ Transform ▾ Discover ▾

Aft Faults: AFT

AFT_FAULTS_EXT...
LOGBOOK

Clear All

REFINE BY

- AFT_FAULTS_EXTENDED
 - accelerator
 - creator
 - elogbook_fault_id
 - fault_description
 - fault_id
 - fault_state
 - full_name
 - logbook_faulty_element
 - op_duration
 - op_end_time
 - parent_fault_id
 - prevents_injection
 - start_time
- Other
 - comment_terms
 - event_comment
 - fault_groupname

Filter...

Cryogenics

QPS

Power converters

SPS

Technical Services

RF

Miscellaneous

Collimator

Beam instrumentation

Injection

Access

Beam dump

PS

Vacuum

Controls

PSB

Experiments

BIS

Operation

SIS

Select All

fault_groupname: AFT

logbook_nickname: LHC OP TESTS

COLUMN CONTAINER

CHART

faultname

- ALICE
- ATLAS
- BCT
- BLM - hardware
- BPM - hardware
- CMS
- Cold compressor
- Controller
- Controls
- Cooling
- CTRL
- CV
- Electrical Services
- Feedback
- Hardware
- Infrastructure
- Interlock
- LHCb
- Mains glitch
- Middleware
- Miscellaneous
- Network
- No beam
- Operator error
- Pressure

CHART

faultname

- Cold compressor
- Hardware
- No beam
- Controls
- VOID
- Electrical Services
- Interlock
- Controller
- Voltage source
- Software
- BLM - hardware
- CV
- Mains glitch
- Cooling
- CTRL
- Infrastructure
- Pressure
- Transverse feedb...
- Miscellaneous
- WorldFIP
- ALICE
- Middleware
- Feedback
- Valve
- BPM - hardware

COLUMN CONTAINER

Drag portlets below to nest them.

RESULTS TABLE

Other 0 RECORDS SELECTED

comment_ter...	event_comment	fault_groupname	faultname	linenames	logbook_nickname	op_r...
confirmation	Global Post Mortem Event Confirmation	RF	Hardware	BeamSetup	LHC OP	BEAM IN
compensat...	Ramping down dipole, solenoid and compensators for Alice's access.	Access	Hardware	BEAM	LHC OP	SETTING_U
LHCb	Took over the machine during squeeze for MD to test separate collision beam process for LHCb (collide IPs 1,2,5 first, then IP8 aft...	SPS	No beam	PROTONPHY	LHC OP	BEAM IN
update	TIM update done... includes update of DIP publication restarts	Vacuum	Valve	BEAM	LHC OP	SETTING_U
	PCC5 on RSF2.A67B2	Miscellaneous	VOID	BeamSetup	LHC OP	NO BEAMS
bpmhlc, ca...	LHC SEQ: BPMLHC calibration finished. Overall result: SUCCESS	Beam dump	Controls	PROTONPHY	LHC OP	SETUP
BOMONT, ...	MARC, BOMONT (MBOMONT) has been assigned RBAC Role: PO-LHC-Piquet and will expire on: 2011-07-12 12:40:00.000000	QPS	Controller	RECOVERY	LHC OP	SETUP
lhc	LHC SEQ: MCS checks finished	BIS	Hardware	PROTONPHY	LHC OP	SETUP
transaction	org.springframework.transaction.TransactionSystemException: Could not roll back JDBC transaction; nested exception is java.sql... With the 6 it is better, carry on with 144 bunches.	Controls	Infrastructure	MD	LHC OP	BEAM IN
		Miscellaneous	VOID	PROTONPHY	LHC OP	BEAM IN
LHCb, finish	LHCb access finish.	RF	Hardware	ProtNucPhy	LHC OP	NO BEAMS
ctrl, lhc	LHC RUN CTRL: BEAM MODE changed to RAMP DOWN	Cryogenics	Cold compressor	PROTONPHY	LHC OP	SETUP
	LHC SEQ: B1 Collimators to parking finished	QPS	Hardware	IONPHY	LHC OP	SETUP
QPS, acces	we switch off sector 45 and 56 for the QPS acces in point 5	QPS	Hardware	PROTONPHY	LHC OP	NO BEAMS

VIEW OPTIONS ▾ ACTIONS ▾

TAG CLOUD

accelerator B1 b1 B2 b2 bct beam beams

bim BLM bpm bpmhc calibration checks CMS

coll collimator collimators confirmation cryo

ctrl dipole dump EPC fill handshake injection

interlock interlocks lbd LHC LHC Piquet

piquet post precycle QPS ramp RBAC RP sanity

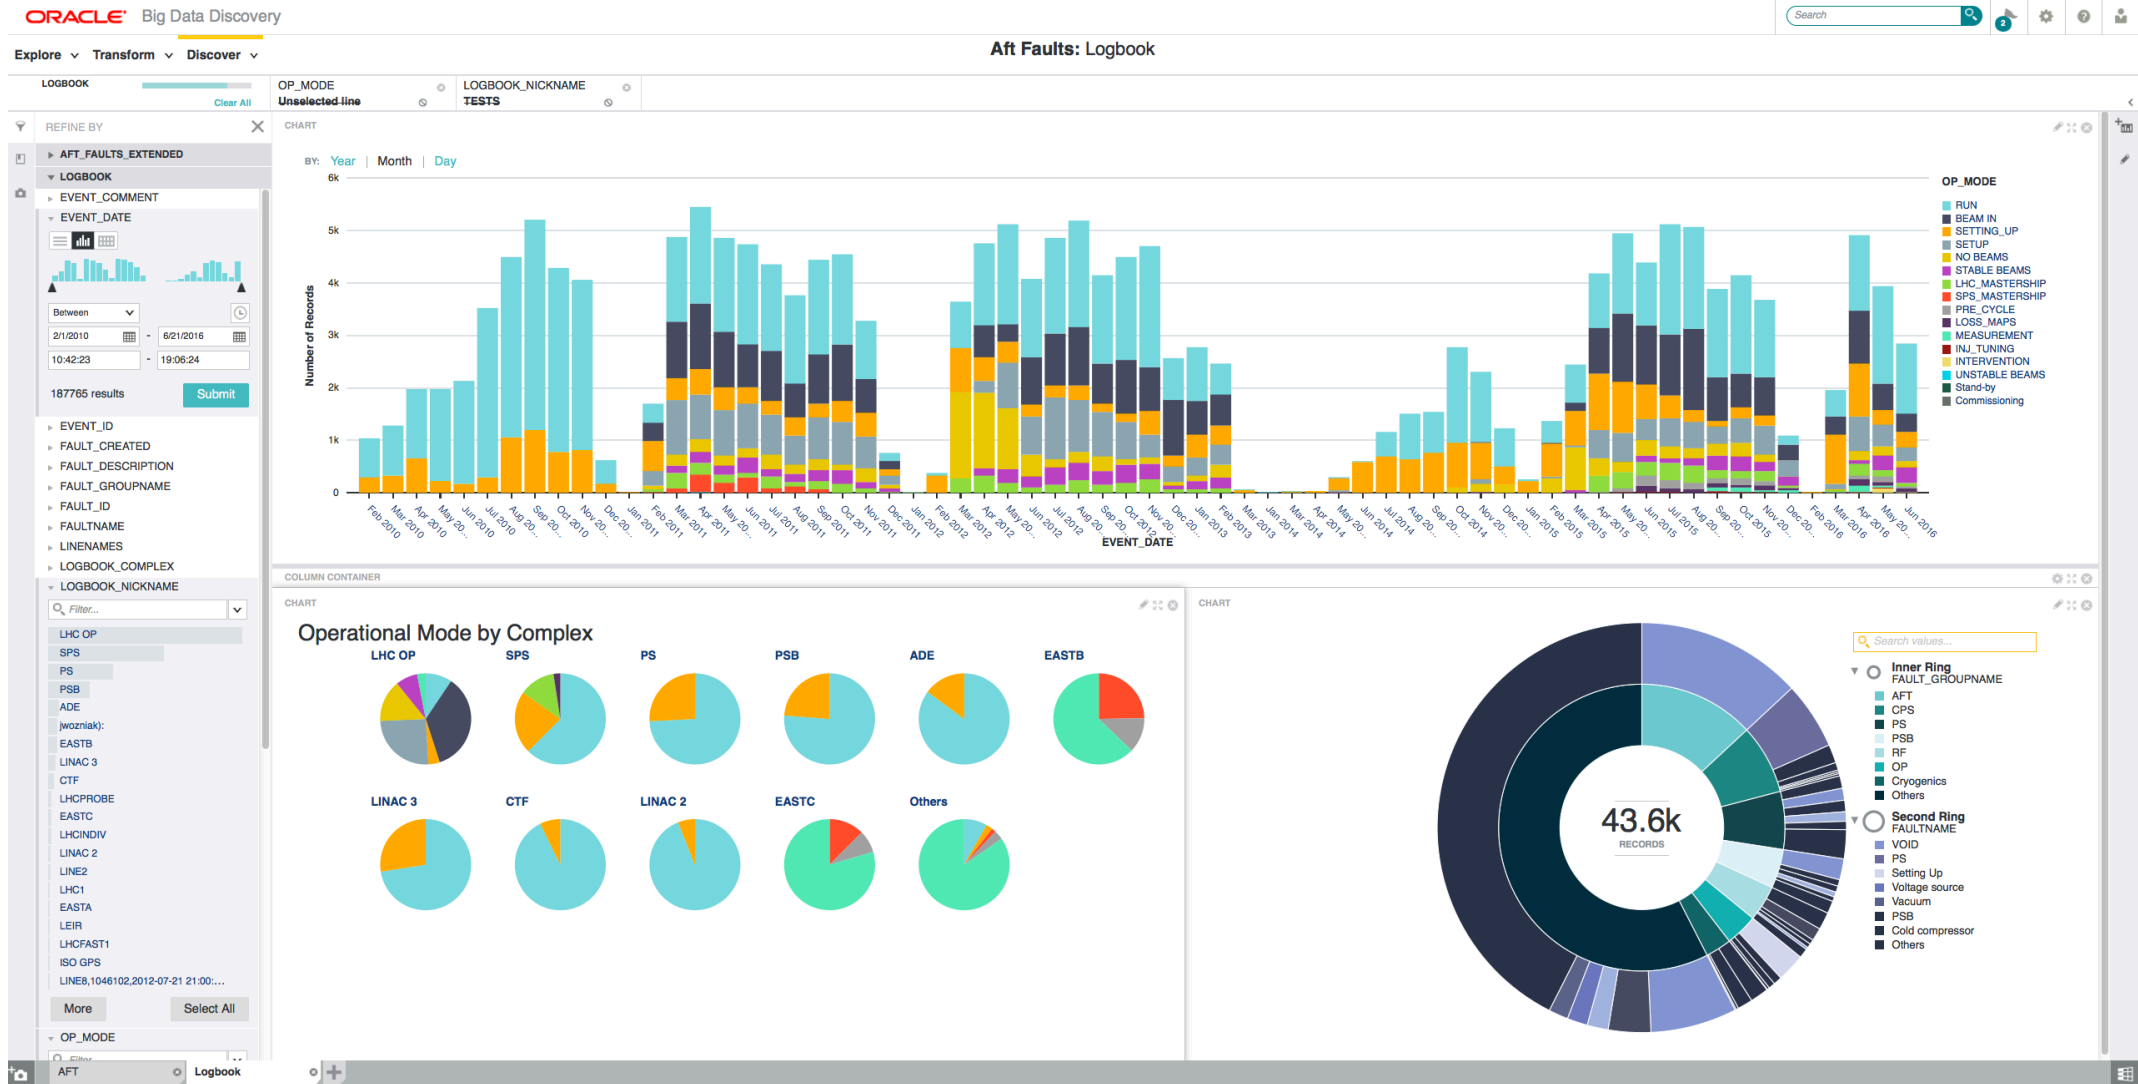
settings setup solenoid tdq thresholds vacuum

valves warnings XPOC

Explore comment_terms by Number of records



Discovery Applications



Discovery Applications



Explore Transform Discover

Aft Faults: AFT

AFT_FAULTS_EXT... Clear All
fault_groupname
Cryogenics

REFINE BY

- AFT_FAULTS_EXTENDED
 - accelerator
 - creator
 - elogbook_fault_id
 - fault_description
 - fault_id
 - fault_state
 - full_name
 - logbook_faulty_element
 - op_duration
 - op_end_time
 - parent_fault_id
 - prevents_injection
 - start_time
 - Other
 - comment_terms
 - event_comment
 - fault_groupname
- Filter...
- AFT
 - QPS
 - Power converters
 - SPS
 - Technical Services
 - RF
 - Miscellaneous
 - Collimator
 - Beam instrumentation
 - Injection
 - Access
 - Beam dump
 - PS
 - Vacuum
 - Controls
 - PSB
 - Experiments
 - BIS
 - Operation
 - SIS
- More Select All
- faultname
 - linenames
 - logbook_nickname
 - op_mode
 - Expand selection
 - LOGBOOK_CLEAN



COLUMN CONTAINER

RESULTS TABLE

Other 0 RECORDS SELECTED

comment_ter...	event_comment	fault_groupname	faultname	linenames
ghh	GHH & Reyes	Cryogenics	Cold compressor	IONPHY LF
compensat...	LHC SEQ: ramping up ALICE DIPOLE and COMPENSATORS	Cryogenics	Cold compressor	ACCESS LF
beams	- started the shift with a stable beams happily ongoing (Fill 2309).	Cryogenics	Cold compressor	IONPHY LF
I4, lhc, ond...	LHC SEQ: ONDULATOR L4 U_RES RESETED	Cryogenics	Cold compressor	ACCESS LF
	For Philippe	Cryogenics	Cold compressor	IONPHY LF
bctfd, lhc	LHC SEQ: BctFd tests finished. Overall result: SUCCESS	Cryogenics	Cold compressor	ACCESS LF
b2, collima...	LHC SEQ: B2 Collimators to parking starting	Cryogenics	Cold compressor	IONPHY LF
3hrs, MKD...	MKD switch replaced, and now synchronization and local energy scans for 3hrs, then if needed, a change of the delay board. Final...	Cryogenics	Cold compressor	IONPHY LF
LDM, LED...	short access for LED on LDM at UX451	Cryogenics	Cold compressor	ACCESS LF
	LHC SEQ: ONDULATOR L4 ON	Cryogenics	Cold compressor	ACCESS LF
I5b1, qps, ...	RQTL12.L5B1 QPS_OK missing	Cryogenics	Cold compressor	PROTONPHY LF
PM65, acces	end of acces in PM65	Cryogenics	Cold compressor	IONPHY LF
beam	LHC RUN CTRL: BEAM MODE changed to BEAM DUMP	Cryogenics	Cold compressor	IONPHY LF
lhc	LHC SEQ: ALICE & LHC READY for T12 SETUP	Cryogenics	Cold compressor	ACCESS LF
lhcop	'History' application has been captured from lhcop	Cryogenics	Cold compressor	ACCESS LF
lhc	LHC SEQ: MCS checks finished	Cryogenics	Cold compressor	PROTONPHY LF
	Cryo compressor restarted, the recalibration was successful.	Cryogenics	Cold compressor	PROTONPHY LF

TAG CLOUD

B1 b1 B2 b2 bct beam

bpmihc calibration checks

collimators compensators c

confirmation cryo Cryo cryog

dipole EPC fill injection Interi

Ihc LHC LHCb ondu

piquet post precycle QPS ramp

sanity settings setup solenoid t

vacuum valves warnings

Explore comment_terms by Number of records

Inside this project

AFT_FAULTS_EXTENDED REFINEME...

comment_terms - Other

QRL

event_comment - Other

...QRL, around midnight start cooling ma...

...QRL that gave a wrong signal (at 10-4)...

...QRL line C affecting DFB MB. Trying t...

...QRLGE_03R8_TT947. It could be that ...

fault_description - General

...QRL

LOGBOOK REFINEMENTS

EVENT_COMMENT - General

...QRLAB_23R7_TT943.POSST) is actu...

...QRL). All pressure are ok.

Leak in QRL sector 4-5 subsector B45.Q...

...QRL SSB Stable : VGP.230.8R4.Q : 3e...

...QRL that gave a wrong signal (at 10-4)...

EVENT_ID - General

...QRL closed the valves. Clean dump

...QRL,Controls,Cryogenics,2012-09-03 ...

...QRL,Controls,Cryogenics,2012-09-03 ...

It looks like a vacuum gauge on the QRL...

...QRL sector 4-5 subsector B45.Q stable;

3 more refinement matches >

FAULT_DESCRIPTION - General

...QRL

LOGBOOK_CLEAN REFINEMENTS

event_comment - General

...QRL that gave a wrong signal (at 10-4)...

...QRL).

...QRL SSB

Leak Sector 4-5 QRL subsector B : stable

...QRL, around midnight start cooling ma...

6 more refinement matches >

fault_description - General

...QRL

Save as default state



Notebooks

- Easy to create and share interactive documents that contain code
- Step by step execution reproduce the analysis, charts, etc.
- Support for multiple languages/kernels
- Multiple notebook software available
 - Jupyter/IPython
 - BDD provides notebook from version 1.2.0 (BDD Shell)
 - Can be used with Jupyter/IPython
 - HUE notebooks
 - Apache Zeppelin
 - More...
- SWAN cloud notebooks at CERN
 - Collaboration to integrate with Hadoop clusters and software



FCC RAMS studies: Cryogenic valves reliability

- Reliability of degrading components of valves of the CERN Large Hadron Collider (LHC) University of Technology
 - Analysis of the difference between request and feedback aperture
 - Coding in Matlab
- Next: Scale processing using Spark
 - Process valves in parallel/distributed
 - Generate charts for each valves
 - Notebooks using PySpark
 - Collaborative review and improvement of the analysis
 - Share results

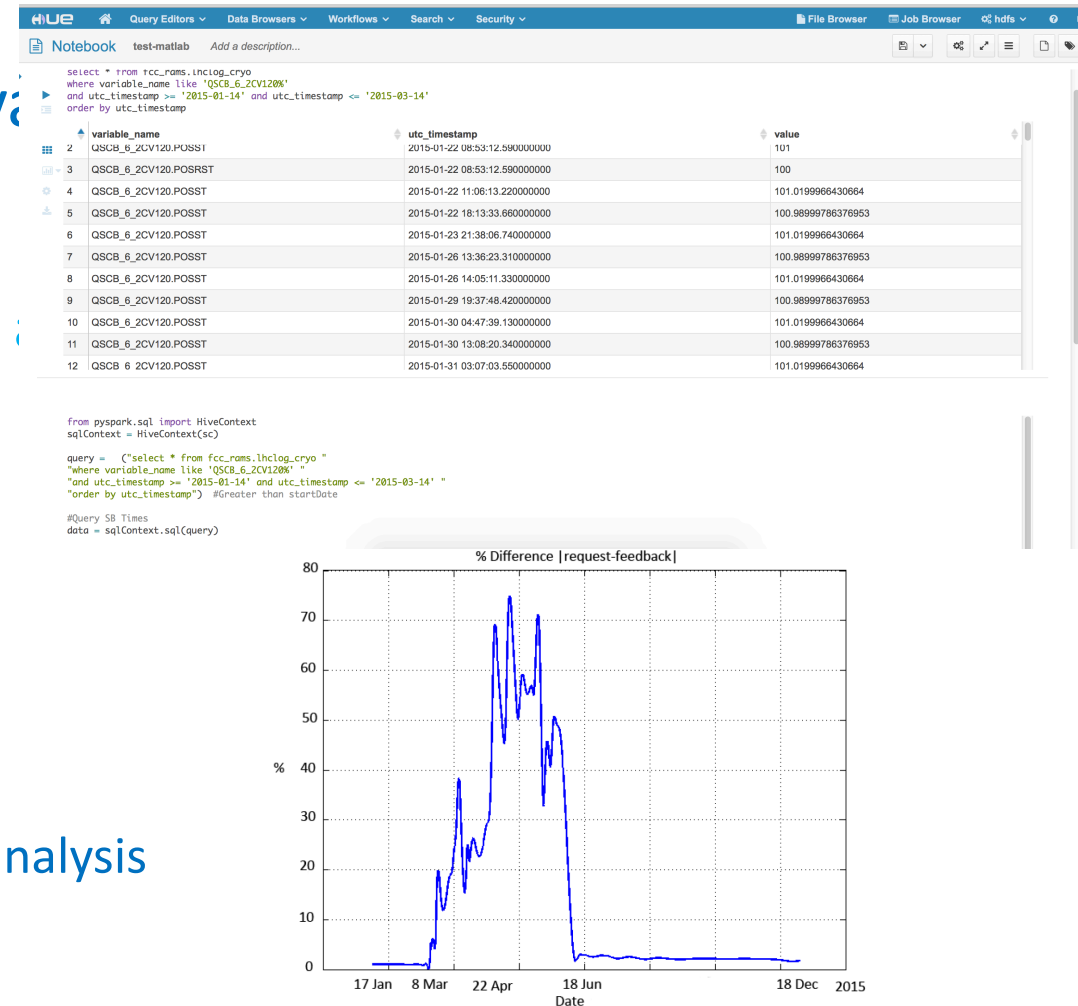


Figure 6. Absolute difference (%) between request and feedback for QSCB_6_2CV120

Conclusions

- Overcome technical limitations for several CERN use cases
 - Unlocks new ways to exploit your data investment
- Great ecosystem for data ingestion, processing, analytics, SQL engines, etc.
- Heterogeneous data access
- Important to facilitate the interaction with data
 - Data visualization and discovery
 - Helps users to browse, explore and understand their data
 - Combined with ETL for data cleaning and feature engineering
 - Notebooks are easy to use and powerful for advanced analytics
 - Self-service tools improve productivity
 - Users should be able to do what they need without IT intervention
 - Collaborative environment



www.cern.ch