



6 May 2014

CERN openlab IT Challenges workshop, <https://indico.cern.ch/event/297559/>

CERN and Data Analytics Challenges


Kacper Szkudlarek, CERN
Manuel Martin Marquez, CERN
Eric Grancher, CERN



Outline

- CERN and openlab
- Analytics challenges
 - Analytics investigation
 - Fault prediction
 - CERN openlab analytics challenges
 - Technology evolutions
- Analytics as a service

Events at LHC

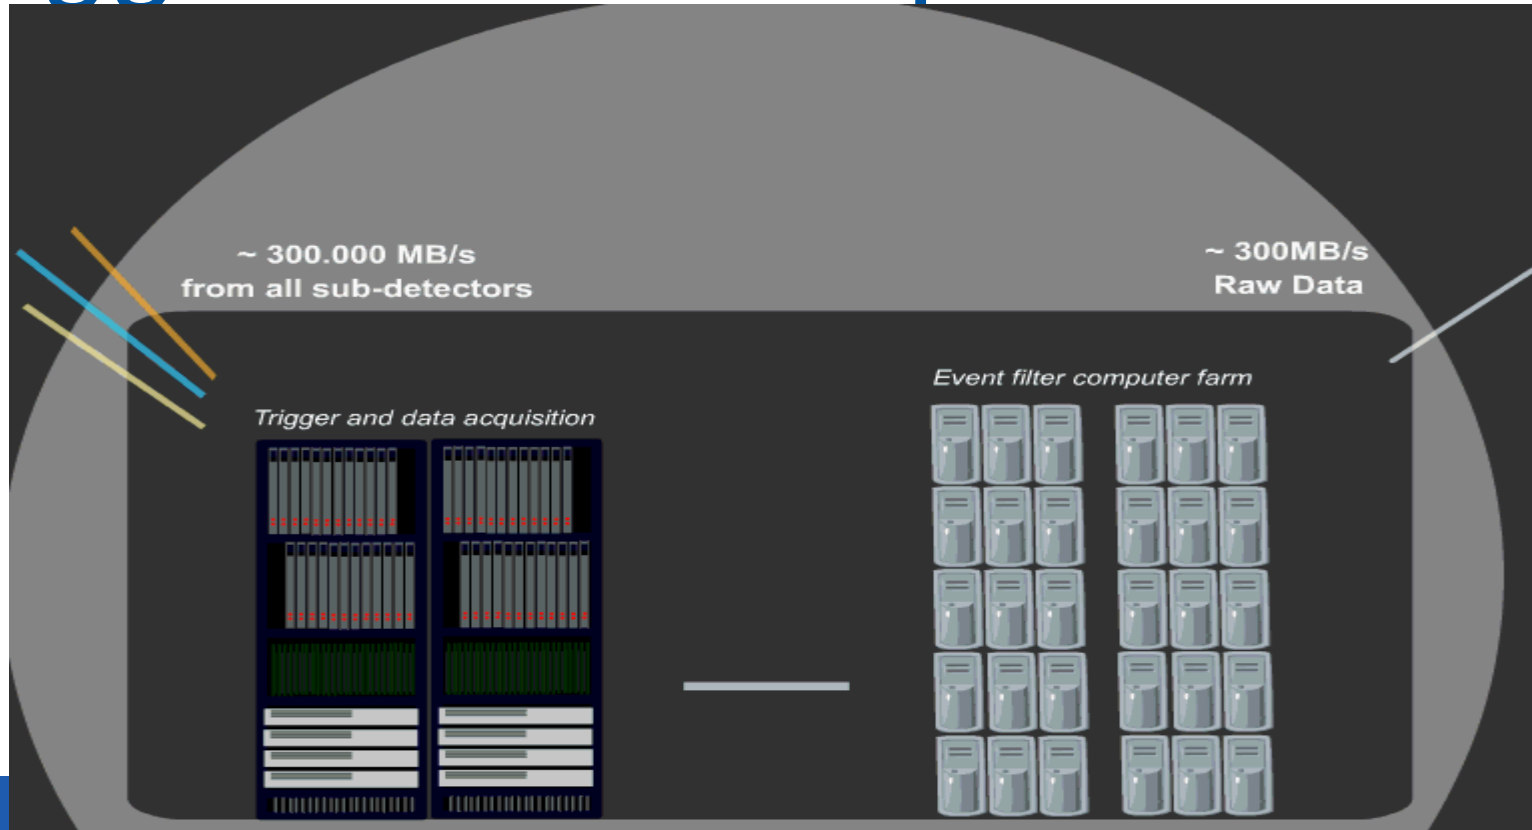


Luminosity :
 $10^{34} \text{cm}^{-2} \text{s}^{-1}$


40 MHz – every 25 ns

20 events overlaying

Trigger & Data Acquisition



Data Recording



of cores: 99,061
of disks: 80,754
of processors: 18,633
of 10GB NICs: 3,218
of 1GB NICs: 19,334
of servers: 10,809
Disk space (TiB): 129,455
RAM memory (TiB): 349



Credit: Alberto Pace

Distributed analysis in
the different sites
(« LHC Computing
Grid »)

CERN




CPU



Network



Databases



Storage



Infrastructure

openlab (1/3)

- Public-private partnership between CERN and leading ICT companies, currently in fourth phase (started in 2003)
- Its mission is to accelerate the development of cutting-edge solutions to be used by the worldwide LHC community
- Innovative ideas aligned between CERN and the partners, for products “you make it, we break it”

Partners



ORACLE

SIEMENS

Contributors



Associates

Yandex

openlab (2/3)

- Many successes (DB competence center/Oracle):
 - RAC on Linux x86 (9.2 PoC, 10.1 production with ASM),
 - Additional required functionality (IEEE numbers, OCCI, instant client, etc.),
 - WinCC OA and RAC scalability,
 - Monitoring with Grid Control,
 - Streams world wide distribution,
 - Active DG, GoldenGate,
 - Analytics for accelerator, experiment and IT,
 - Etc.

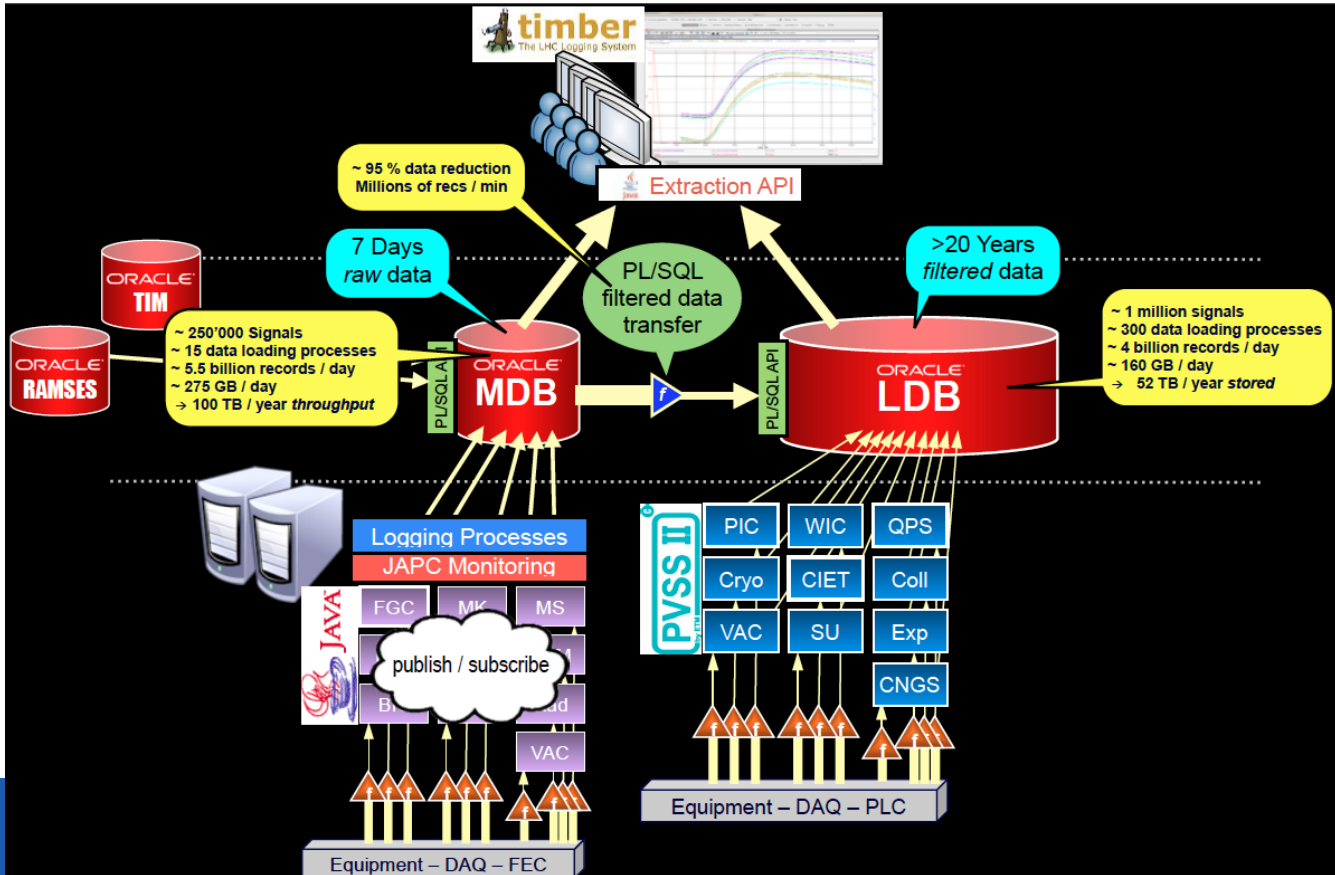


Networked devices

- Large number of control devices, front-end equipment, etc.
- Many critical systems: Cryogenics, Vacuums, Machine Protection, etc.
- (Also called “Internet of Things”)
- Example: LHC logging service (BE-CO)

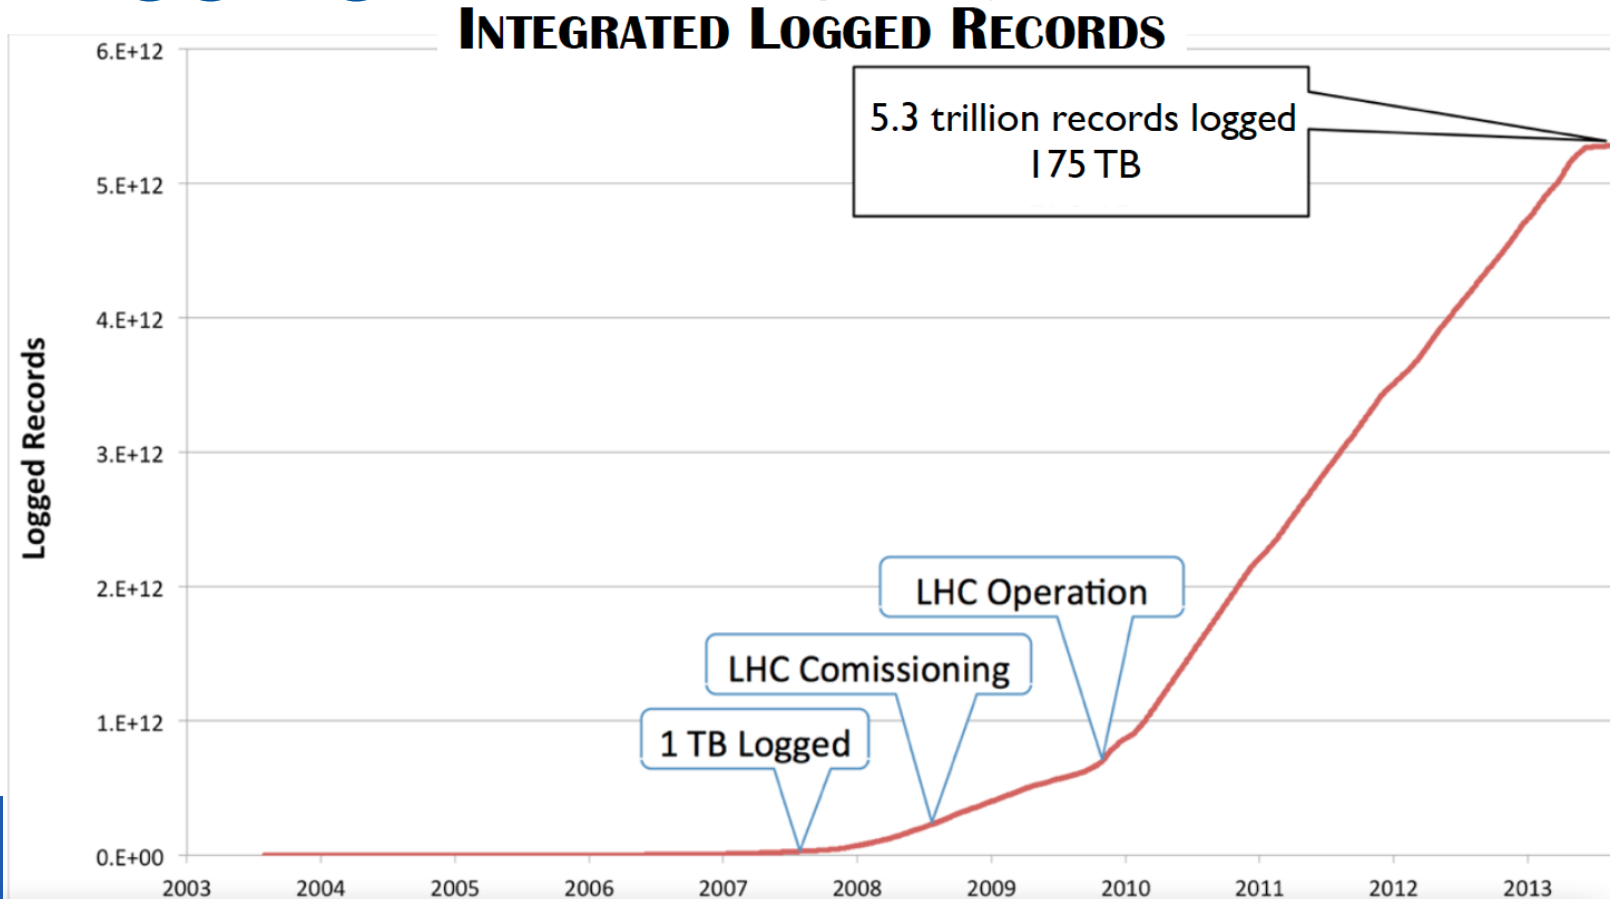
Logging service (1/3)

Credit: Chris Roderick




Logging service (2/3)

Credit: Chris Roderick



Logging service (3/3)

Credit: Chris Roderick



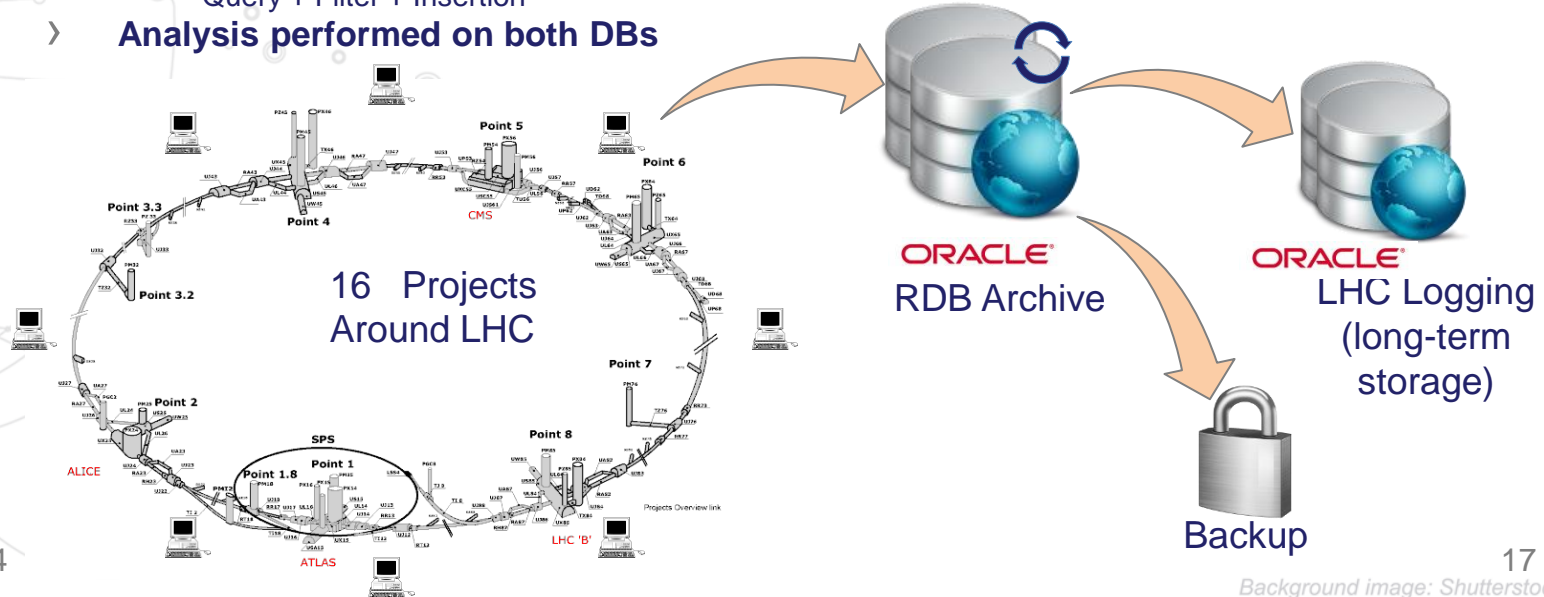


Improvements for Archiver of WinCC Open Architecture

Report on work by departments:
Information Technology
Beams
Engineering

Use case: Quench Protection System

- › **Critical system for LHC operation**
 - Major upgrade for LHC Run 2 (2015-2018)
- › **High throughput for data storage requirement**
 - Constant load of 150k changes/s from 100k signals
- › **Whole data set is transferred to long-term storage DB**
 - Query + Filter + Insertion
- › **Analysis performed on both DBs**



Performance Optimizations

- › **Use of Oracle Index Organized Tables**
- › **Tuning of DB data queries**
 - Search predicates, time-based partitioning
 - Alignment of data in RAC cluster
- › **Changes in database schema**
 - Focus on redo log and space reduction
- › **Tuning of database parameters**

BE-CO


IT-DB

EN-ICE

Results summary

- › **Nominal conditions**
 - Stable **constant load of 150k changes/s**
 - 100 MB/s of I/O operations
 - 500 GB of data stored each day
- › **Peak performance**
 - Exceeded **1 million value changes per second**
 - 500-600 MB/s of I/O operations
- › **All CERN production WinCC OA systems (accelerators, detectors and technical infrastructure, 600 servers) will benefit from these optimizations**
- › **Next challenge: ~10x increase**
 - Required for next major upgrade (2019-2020)

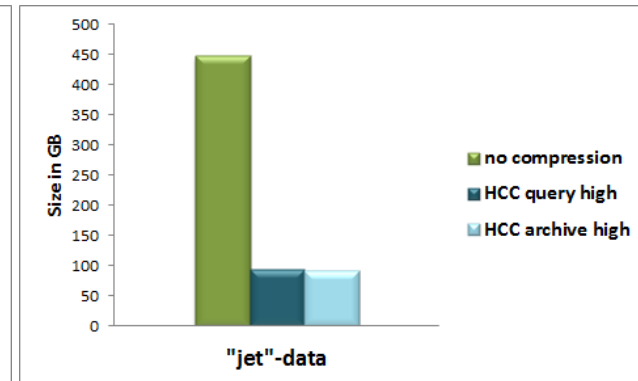
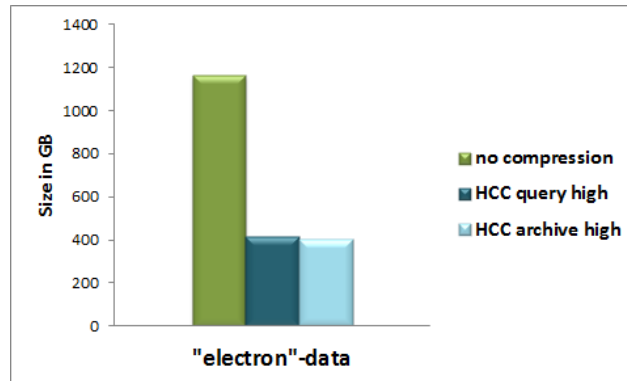
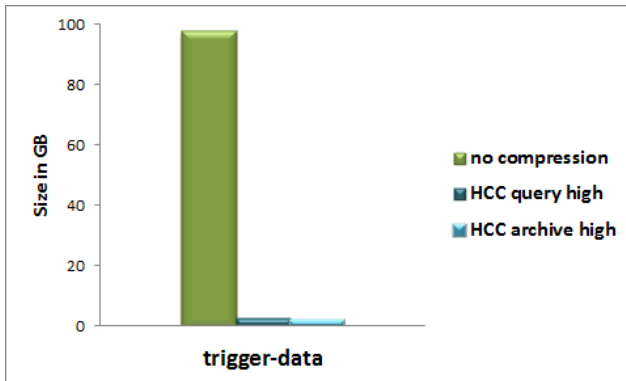
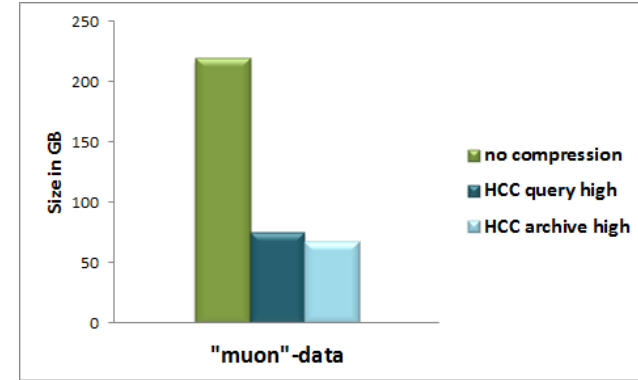
I/O Megabytes per Second



Technology evolutions (1/4)

Credit: Maaike Limper

- Columnar compression



Technology evolutions (2/4)

- In-memory databases
 - Commodity servers with 6TB memory or more
 - Together with compression, enable to have (for some applications) the active data set fully in memory
- Flash with low latency storage

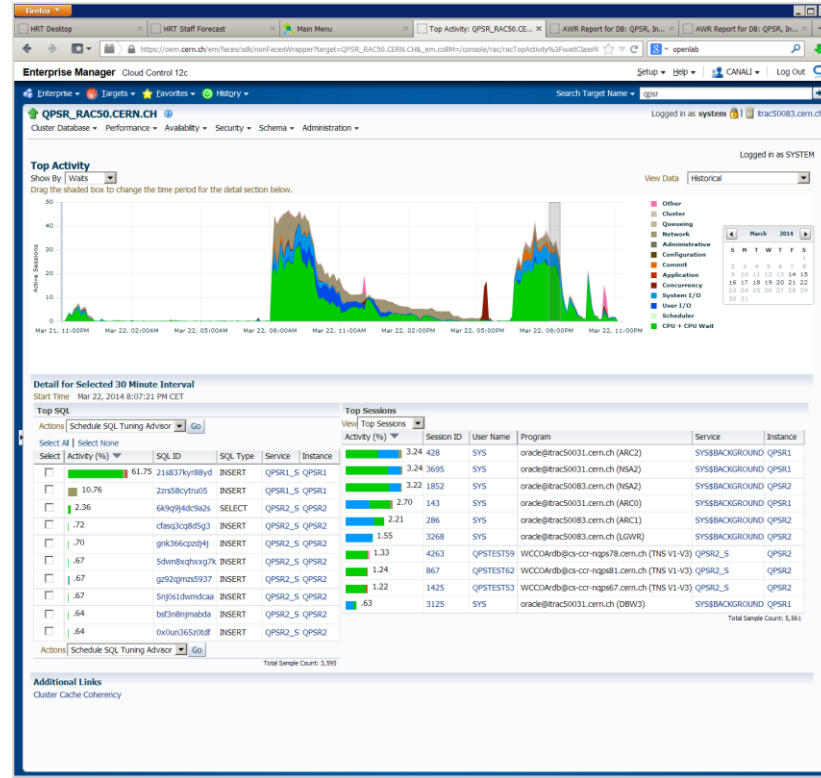
Technology evolutions (3/4)

- Networking
 - 10GbE and more
 - (security...)

Technology innovation (4/4)

Credit: Luca Canali

- Scalable database and storage infrastructure
- Powerful analysis packages like ROOT and R



Anomaly detection


SVM - Support Vector Machines

- 1) Build a **SVM** (*Neural Network like*) **model**
 - **self trained**
 - no supervision

- 2) Predict and compare:

Real data Vs Prediction

- 1) **Blindly recognize anomalies**
- 2) **No other information required** (*i.e. thresholds*)




Information discovery

URL link

Manuel Martin Marquez, CERN IT-DB

Experiment data placement

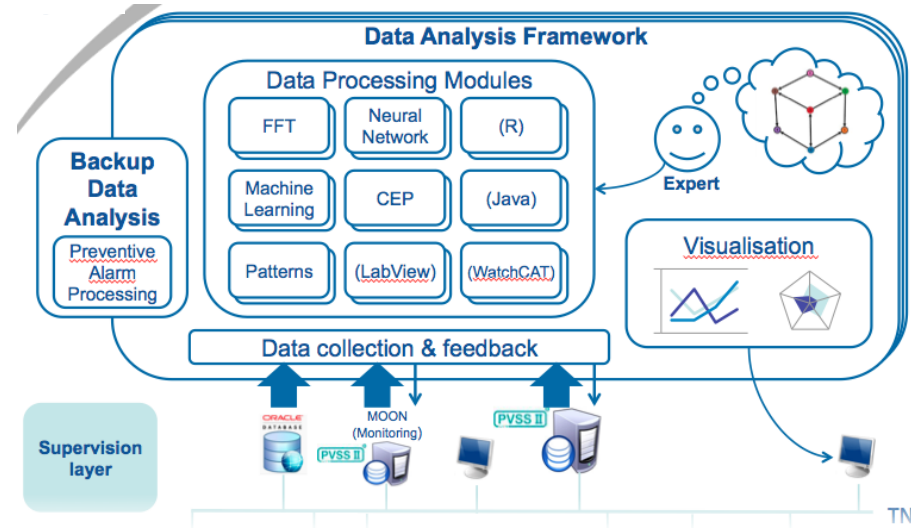
- Intelligent data placement models for the CMS and ATLAS experiments
- Need to extract further knowledge from the monitoring data in order to implement an effective data placement
 - Correlate file-access monitoring with site status
 - Readiness, queue length, storage and CPU available
 - Classify analysis activities and needed resources
 - Making recommendations
 - Learn from the past trends and patterns



Analytics as a service

- “Analytics platform” or (Big data) “Analytics-as-a-service” (A³S ?):
- Data fed from multiple sources (live)
- Stored reliably
- Data processing with multiple systems
- Easy access, domain expert natural language (DSL)
- Visualisation

Credit: CERN EN-ICE



TN

Conclusion

- CERN: very diverse and challenging requirements
- CERN openlab IV, a lot done on Analytics
- CERN openlab V, a lot to be done, CERN and others (EMBL-EBI, ESA, etc.), whitepaper)
<http://cern.ch/openlab/>
- “you make it, we break it”
- Complex integration challenges of rapidly evolving technologies
- Some more about Oracle and CERN... DB12c



www.cern.ch