

CERN


European Organization for Nuclear Research  
Organisation Européenne pour la Recherche Nucléaire

# Big Science Meets Big Data

Dr Bob Jones

CERN

Bob.Jones <at> CERN.ch



# CERN was founded 1954: 12 European States Today: 20 Member States



~ 2300 staff  
~ 790 other paid personnel  
> 10000 users  
Budget ~1000 MCHF annual

**20 Member States:** Austria, Belgium, Bulgaria, the Czech Republic, Denmark, Finland, France, Germany, Greece, Hungary, Italy, Netherlands, Norway, Poland, Portugal, Slovakia, Spain, Sweden, Switzerland and the United Kingdom

**Candidates for Accession:** Romania, Israel

**Observers to Council:** India, Japan, the Russian Federation, the United States of America, Turkey, the European Commission and UNESCO

Video: <http://cds.cern.ch/record/1541893>

Data flow to permanent storage: 4-6 GB/sec

CERN Computer Centre


LHCb ~ 200-400 MB/sec

ATLAS ~ 1-2 GB/sec

ALICE ~ 1.25 GB/sec

CMS ~ 1-2 GB/sec


Bob Jones - December 2012





# Big Data !

IT



Experiment: 7000 tons  
150 million sensors  
Generating data 40 million times per second

→ Peta Bytes / sec !



Trigger  
Select 100,000 per second

→ Tera Bytes / sec !



HLT / Filter  
Select 100 per second

→ Giga Bytes / sec !

# Tier-0: Central Data Management

- **Hierarchical Storage Management: CASTOR**
  - Rich set of features:
    - Tape pools, disk pools, service classes, instances, file classes, file replication, scheduled transfers (etc.)
  - DB-centric architecture
- **Disk-only storage system: EOS**
  - Easy-to-use, stand-alone, disk-only for user and group data with in-memory namespace
    - Low latency (few ms for read/write open)
    - Focusing on end-user analysis with chaotic access
    - Adopting ideas from other modern file systems
    - Running on low-cost hardware (JBOD and SW RAID )

# Active tapes


- Inside a huge storage hierarchy tapes may be advantageous!



We use tape storage products from multiple vendors

# CASTOR current status


(June 2013)



77 petabytes across 362 million files




# Data Handling and Computation for Physics Analysis



Online trigger and filtering

Selection & reconstruction


Offline Reconstruction



Processed Data (Active tapes)



100% Raw data



Event reprocessing


Batch physics analysis

1%

Offline Analysis w/ROOT



Analysis objects (extracted by physics topic)



Interactive analysis


Event simulation

Offline Simulation w/GEANT4



# WLCG – what and why?

- A distributed computing infrastructure to provide the production and analysis environments for the LHC experiments
- Managed and operated by a worldwide collaboration between the experiments and the participating computer centres
- The resources are distributed – for funding and sociological reasons
- Our task was to make use of the resources available to us – no matter where they are located
- Secure access via X509 certificates issued by network of national authorities - International Grid Trust Federation (IGTF)



## Tier-0 (CERN):

- Data recording
- Initial data reconstruction
- Data distribution

## Tier-1 (11 centres):

- Permanent storage
- Re-processing
- Analysis

## Tier-2 (~130 centres):

- Simulation
- End-user analysis

# Broader Impact of the LHC Computing Grid




- WLCG has been leveraged on both sides of the Atlantic, to benefit the wider scientific community
  - Europe (EC FP7):
    - Enabling Grids for E-science (EGEE) 2004-2010
    - European Grid Infrastructure (EGI) 2010--
  - USA (NSF):
    - Open Science Grid (OSG) 2006-2012 (+ extension?)
- Many scientific applications →

Archeology  
Astronomy  
Astrophysics  
Civil Protection  
Comp. Chemistry  
Earth Sciences  
Finance  
Fusion  
Geophysics  
High Energy  
Physics  
Life Sciences  
Multimedia  
Material Sciences

...

# Big Data Analytics

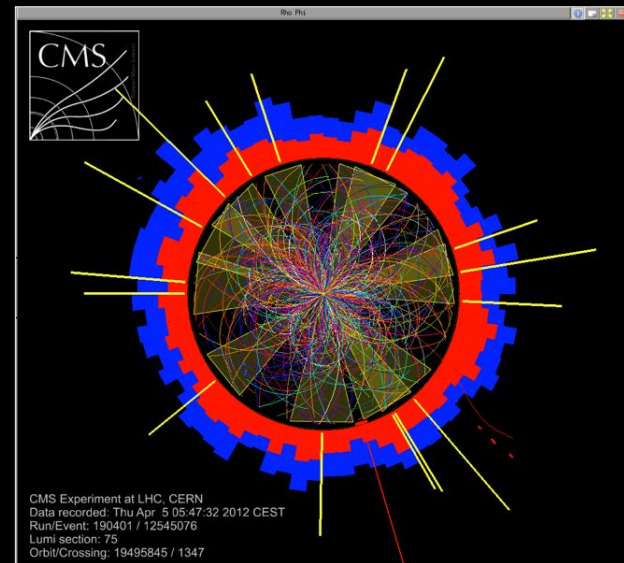
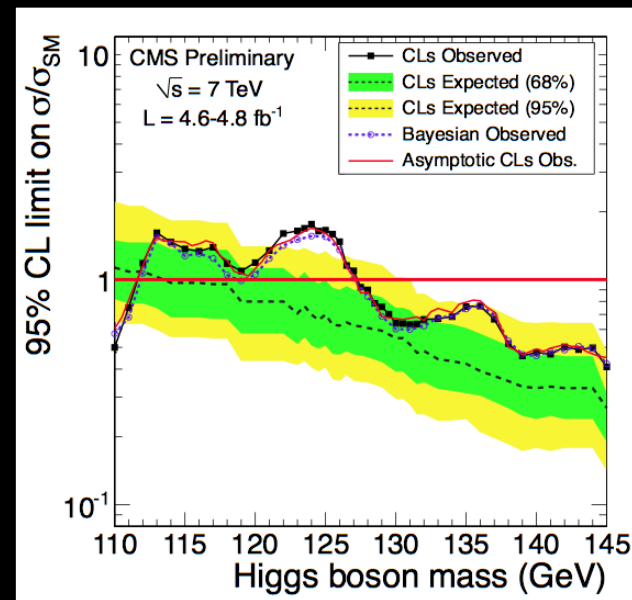
- Huge quantity of data collected, but most of events are simply reflecting well-known physics processes
  - New physics effects expected in a tiny fraction of the total events:
    - “The needle in the haystack”
- Crucial to have a good discrimination between interesting events and the rest, i.e. different species
  - Complex data analysis techniques play a crucial role







# ROOT Object-Oriented toolkit

- **Data Analysis toolkit**
  - Written in C++ (millions of lines)
  - Open source
  - Integrated interpreter
  - Multiple file formats
  - I/O handling, graphics, plotting, math, histogram binning, event display, geometric navigation
  - Powerful fitting (RooFit) and statistical (RooStats) packages on top
  - In use by all our collaborations



# Roofit/RooStats

- Standard tool for producing physics results at LHC
  - Parameter estimation (fitting)
  - Interval estimation (e.g limit results for new particle searches)
  - Discovery significance (quantifying excess of events)
- Implementation of several statistical methods (Bayesian, Frequentist, Asymptotic)
- New tools added for model creation and combinations



# ROOT files

- **Default format for all our data**
- **Organised as Trees with Branches**
  - **Sophisticated formatting for optimal analysis of data**
    - **Parallelism, prefetching and caching**
    - **Compression, splitting and merging**



Over 100 PB stored in this format (All over the world)



# CERN openlab in a nutshell

- A science – industry partnership to drive R&D and innovation with over a decade of success
- Evaluate state-of-the-art technologies in a challenging environment and improve them
- Test in a research environment today what will be used in many business sectors tomorrow
- Train next generation of engineers/employees
- Disseminate results and outreach to new audiences

## PARTNERS



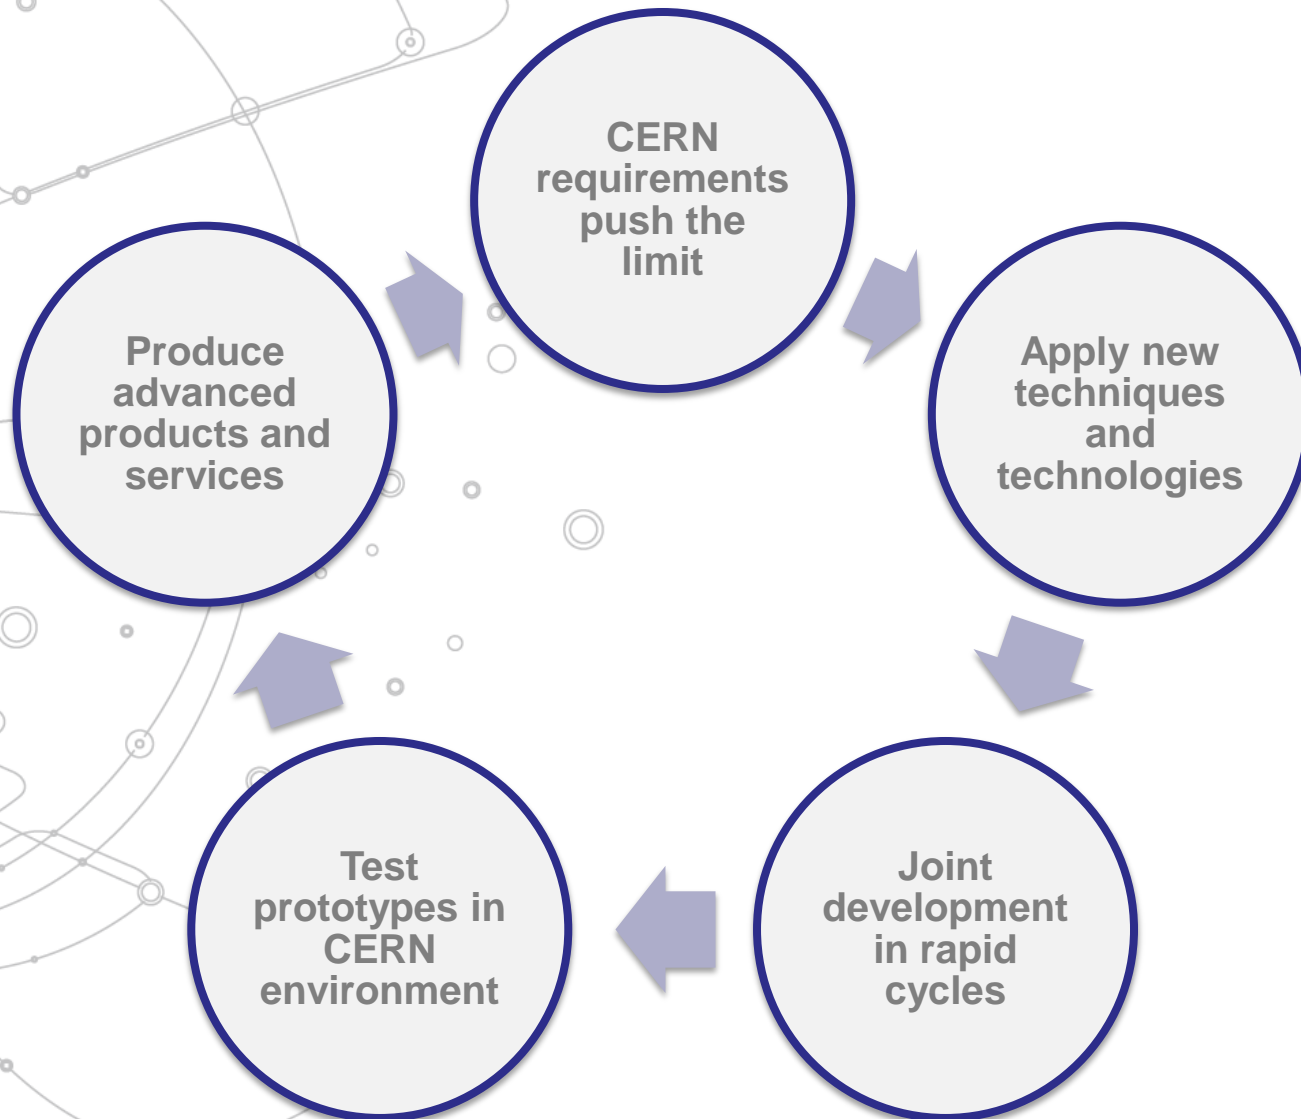
ORACLE

SIEMENS

## ASSOCIATE

Yandex

# Virtuous Cycle

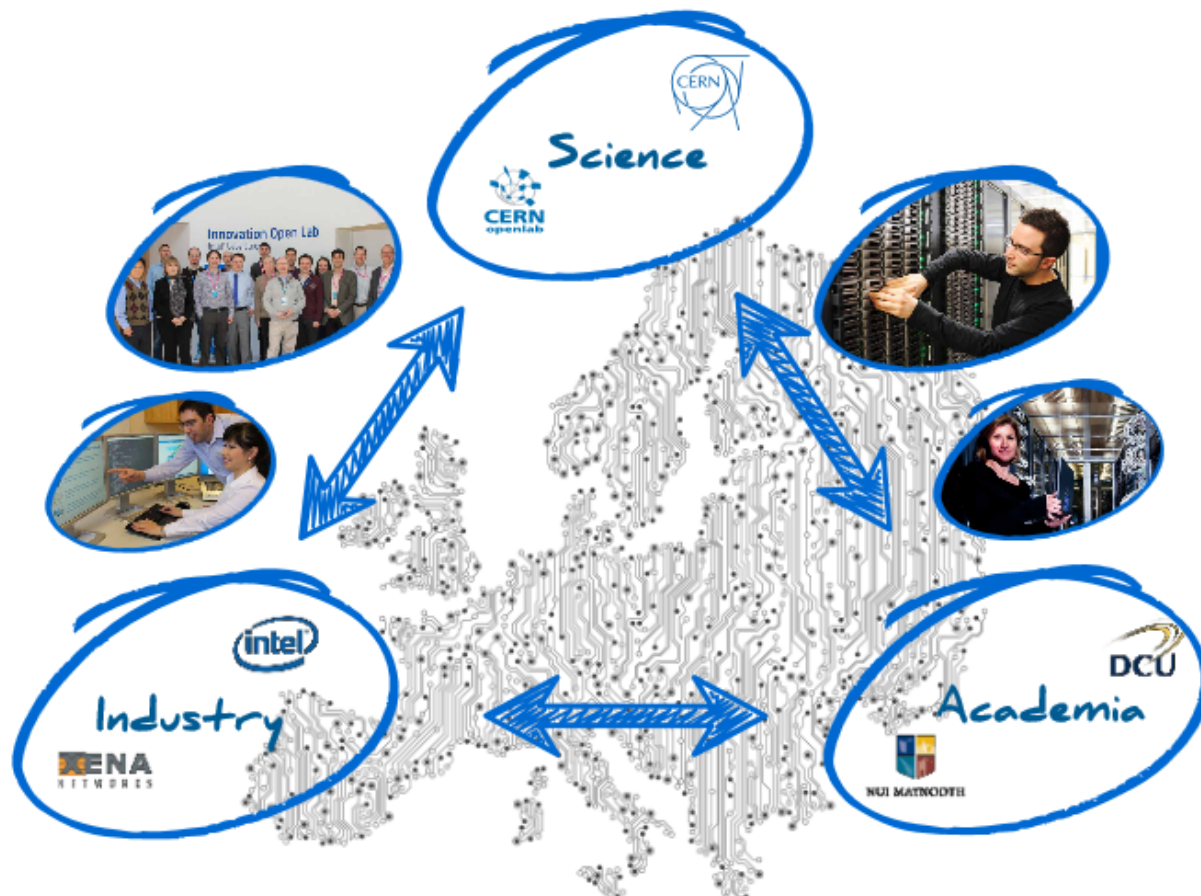


**A public-private partnership between the research community and industry**



# ICE-DIP 2013-2017: The Intel-CERN European Doctorate Industrial Program

» A public-private partnership to research solutions for next generation data acquisition networks, offering research training to five Early Stage Researchers in ICT



Research topics:

- ▶ Silicon photonics systems
- ▶ High speed configurable logic
- ▶ Next generation data acquisition networks
- ▶ Computing solutions for high performance data filtering





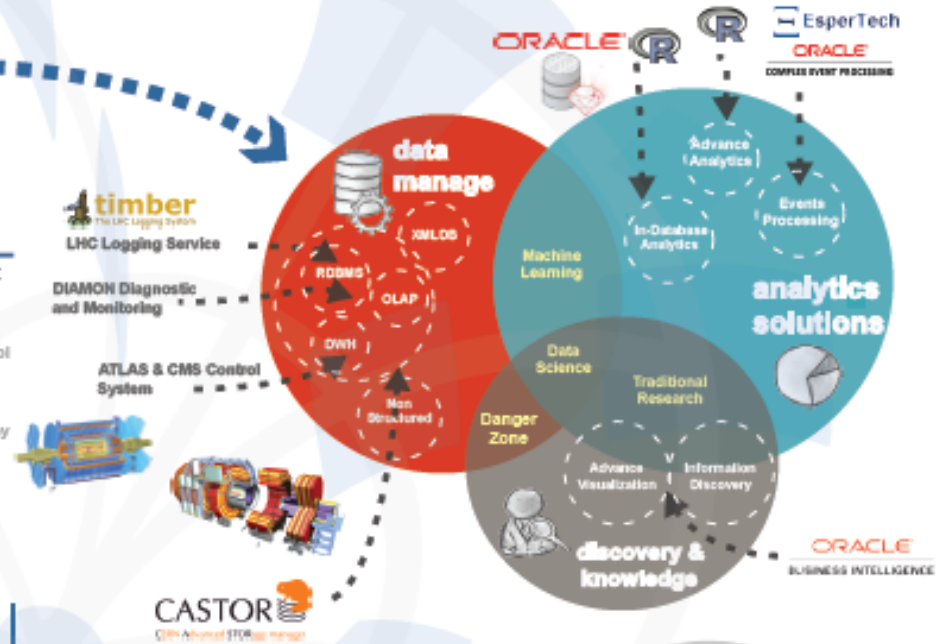
**The openlab Data Analytics Project**

During the past decades, CERN has been gathering and storing an enormous amount of data. This process is costly in terms of technical and human resources. However, the exploitation of the collected data, in other words, the extraction of potential benefits from our data investments, has been pushed into the background or placed on the bottom of our priorities. **Data is the new soil** and it requires nurturing, enriching and managing. This requires additional efforts, but it is clear that those efforts will **generate important value**. The openlab data analytics project aims to change the situation and demonstrates that **small investments in data analytics can lead to big benefits**.




**Data Science and CERN Analytic Environment**

- Beams Department:** Diagnostic and Monitoring system (DIAMON). Status and key performance indicators analysis for the LHC control infrastructure.
- Information Technology:** CERN Advanced STORAGE Manager (CASTOR) Forecasting system evolution for preventing potential incidents by spotting trends and patterns in historical data.
- CMS Experiment:** Automated high voltage trip recovery. Analysis of Complex Event Processing (CEP) conclusions.
- ATLAS Experiment:** Real time monitoring and anomaly detection. Analysis of CEP output data.
- LHC Logging Service:** Oracle R Enterprise for in-database analytics.



**Using Oracle advanced analytics to forecast CERN Accelerators power consumption**

Power consumption forecast has become one of the major research fields in electrical engineering. Forecasting the CERN accelerators complex power consumption is **extremely challenging** due to the **high complexity** and variety of the underlying processes but also **critical given the high level** of consumption, around 1000 GWh. The CERN Electrical Network Supervisor acquires thousands measures per minute for all the subsystems that conform the CERN accelerators complex. Later it stores them in the **CERN Accelerator Logging Service** where the data is persisted in Oracle RAC. Within this context, we demonstrate the importance of performing in-database analytics and how **Oracle Advanced Analytics** has made that possible.



# A European Cloud Computing Partnership: big science teams up with big business



## Strategic Plan

- ▶ Establish multi-tenant, multi-provider cloud infrastructure
- ▶ Identify and adopt policies for trust, security and privacy
- ▶ Create governance structure
- ▶ Define funding schemes



To support the computing capacity needs for the ATLAS experiment



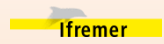
Setting up a new service to simplify analysis of large genomes, for a deeper insight into evolution and biodiversity



To create an Earth Observation platform, focusing on earthquake and volcano research



Adopters




# Initial flagship use cases

## ATLAS High Energy Physics Cloud Use



To support the computing capacity needs for the ATLAS experiment

## Genomic Assembly in the Cloud



A new service to simplify large scale genome analysis; for a deeper insight into evolution and biodiversity

## SuperSites Exploitation Platform



To create an Earth Observation platform, focusing on earthquake and volcano research

- **Scientific challenges with societal impact**
- **Sponsored by user organisations**
- ***Stretch* what is possible with the cloud today**



- Big Data Management and Analytics require a solid organisational structure at all levels
- Corporate culture: our community started preparing more than a decade before real physics data arrived
- Now, the LHC data processing is under control but data rates will continue to increase (dramatically) for years to come: **Big Data at the Exabyte scale**
- CERN is working with other science organisations and leading IT companies to address science's big-data needs



# *Thank you for your attention*



*Accelerating Science and Innovation*