


Computing at CERN



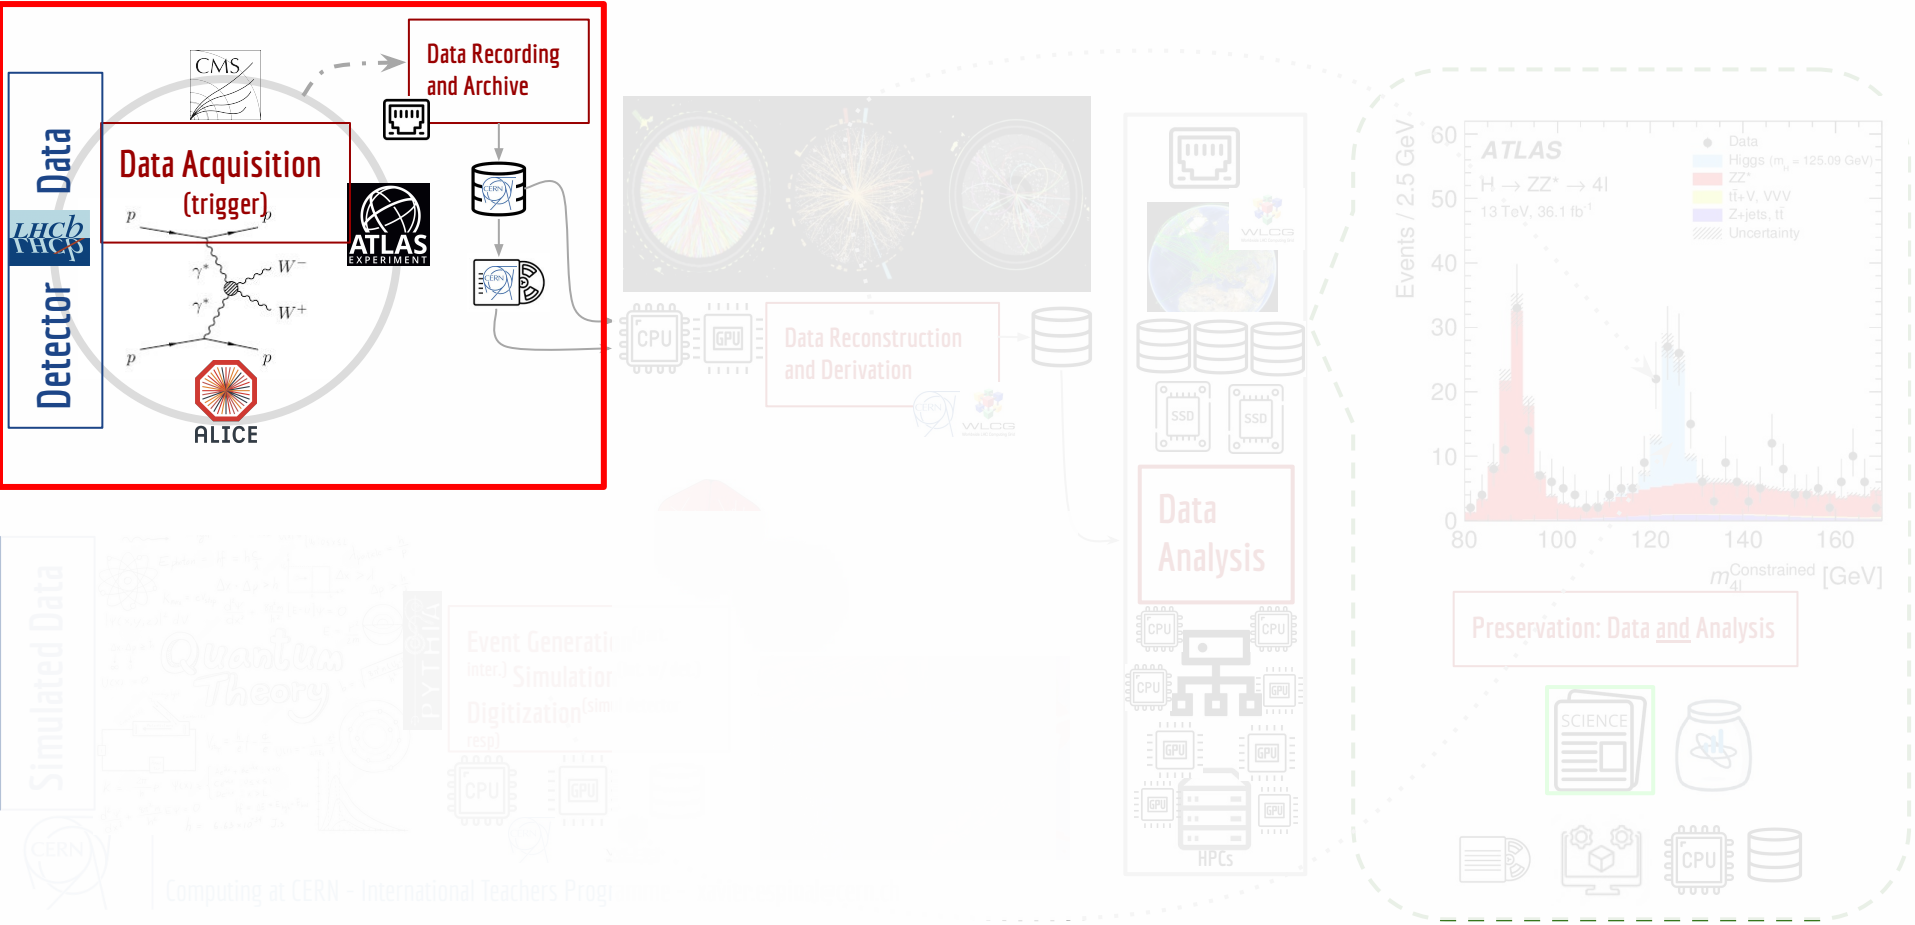
Xavier Espinal (CERN)



Computing at CERN: the big picture



Computing at CERN: the big picture



Our Data *Generators*

Experimental physics is about **DATA!**


CERN experiments transfer all data to the
Computing Center

Data is stored and ready to be processed,
analyzed and shared






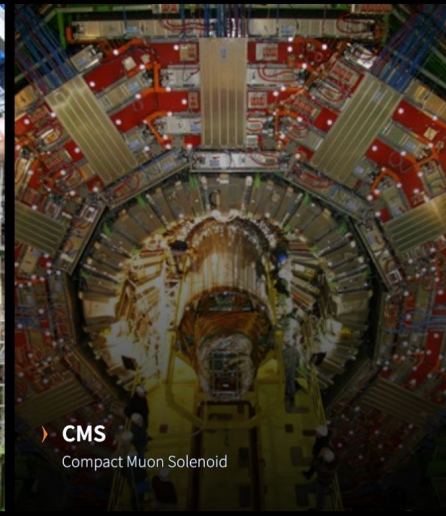
LHC Experiments: the *big 4*




heavy-ion physics, strongly interacting matter at extreme energy densities (quark-gluon plasma)



General-purpose: Higgs boson, extra dimensions, dark matter



General-purpose: Higgs boson, extra dimensions, dark matter



CP violation in the interactions of b-hadrons: matter-antimatter asymmetry of the Universe.

A computing perspective: 0(10-100) Pb/year, 0.5M running jobs (continuous), fat and long-haul networks





LHC Experiments

A computing perspective

Precise measurement of the proton-proton interaction cross section

› **TOTEM**

Total, elastic and diffractive cross-section measurement

Directly search for highly ionizing avatars of new physics that include not only magnetic monopole

› **MoEDAL-MAPP**

Monopole and Exotics Detector at the LHC



Precise measurements of the production spectra relative to neutral particle produced by high energy proton-ion collisions in the very forward region

› **LHCf**

Large Hadron Collider forward

Search for new weakly interacting particles: dark photons, axion-like particles and sterile neutrinos

› **FASER**

Forward Search Experiment

A computing perspective: < Pb/year, data preservation, availability, shareability





The *fixed target* Experiments

Hadron structure and hadron spectroscopy with high intensity muon and hadron beams.

› **COMPASS**
Common Muon and Proton Apparatus for Structure and Spectroscopy

Properties of the production of hadrons in collisions of beam particles (pions, and protons, beryllium, argon and xenon)

› **NA61/SHINE**
SPS Heavy Ion and Neutrino Experiment

Study rare kaon decays. Check some of the predictions the Standard Model makes about short-distance interactions

› **NA62**
North area experiment 62

Radiation process in strong electromagnetic fields

› **NA63**
North area experiment 63

Search for unknown particles from a hypothetical "dark sector", e.g. dark photons, which would carry a new force between visible matter and dark matter

› **NA64**
North area experiment 64

Measure the electric and the magnetic moments of short-lived baryons in a high-energy hadron collider,

› **UA9**
Crystal

› **CLOUD**
Cosmics Leaving Outdoor Droplets

Understand the influence of galactic cosmic rays (GCRs) on aerosols and clouds, and their implications for climate.






A computing perspective: from TeV/y to PB/y, small to large data processing needs, integration of external resources (HPC), workflow synchronisation turnarounds (data quality monitoring)



Antimatter Experiments

A computing perspective



A computing perspective: modest computing requirements, challenge in preservation/reproducibility, data accessibility/shareability





Non-Accelerator Experiments

A computing perspective

Properties of the cosmic rays , origin of dark matter, antimatter and cosmic rays as well as to explore new phenomena.

› **AMS**

Alpha Magnetic Spectrometer



› **CAST**

CERN Axion Solar Telescope

› **OSQAR**

Optical Search for QED Vacuum Bifringence, Axions and Photon Regeneration

A computing perspective: data orchestration, consolidation and custody, from TeV/y to PB/y, some large data processing needs, international collaborations, workflow synchronisation turnarounds (data quality monitoring)





Experimental Facilities

A computing perspective



▶ **ISOLDE**

Isotope mass Separator On-Line facility

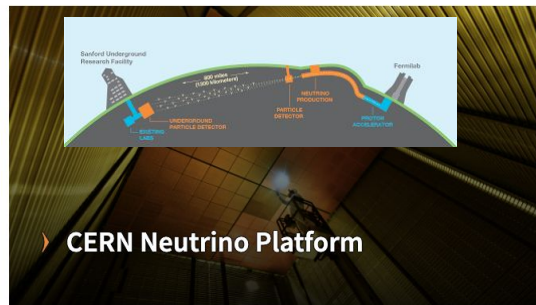


▶ **MEDICIS**



▶ **n_TOF**

Neutron time-of-flight facility

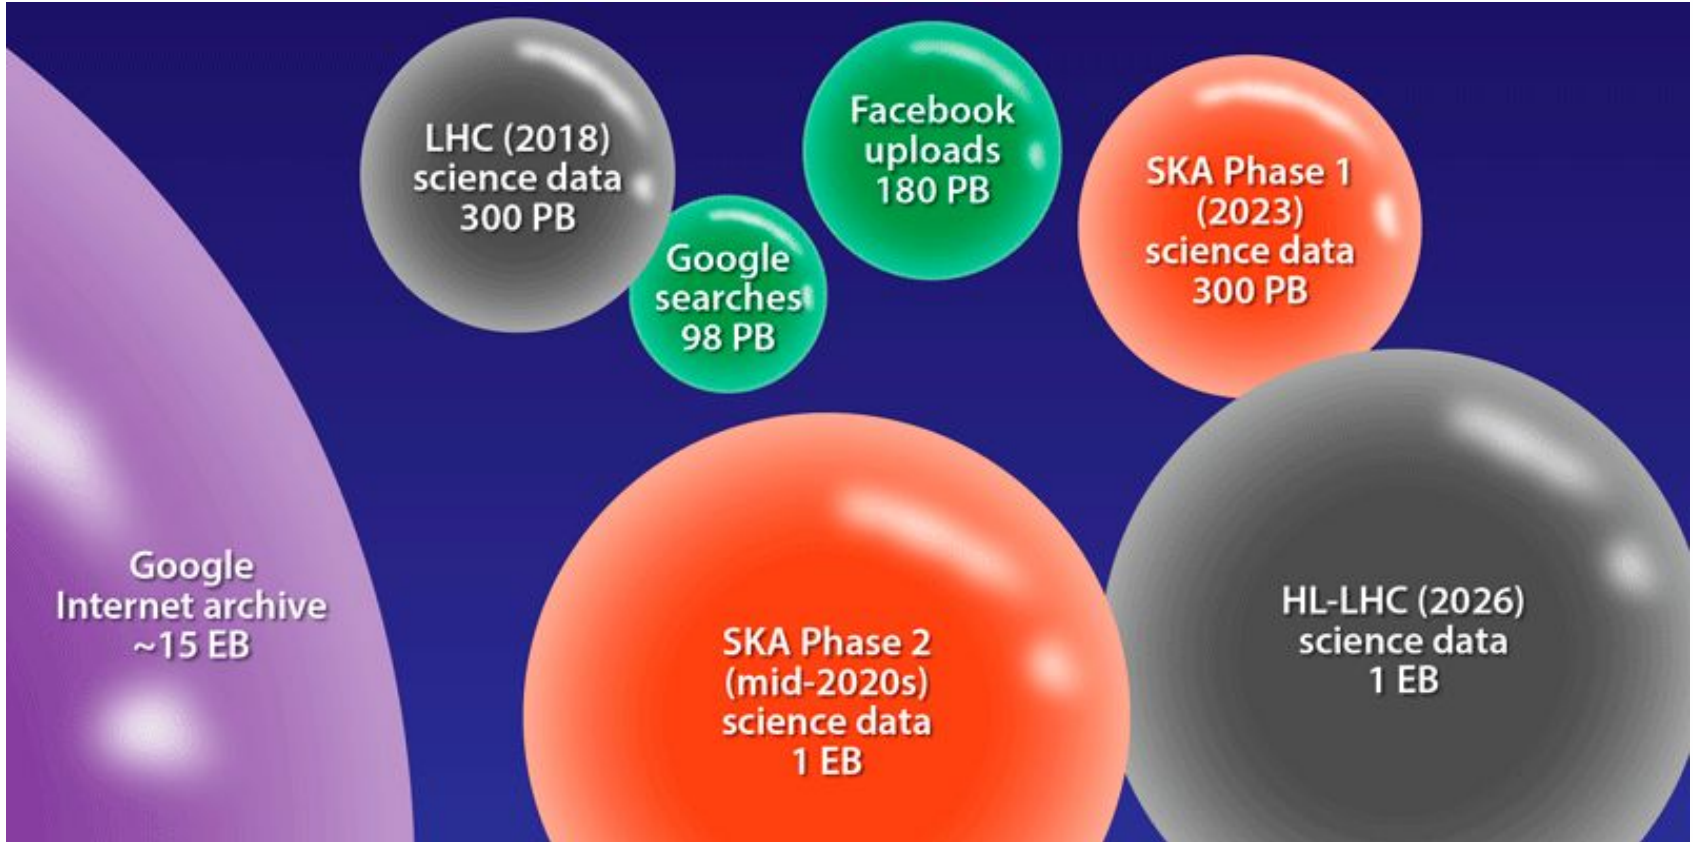


▶ **CERN Neutrino Platform**

A computing perspective: from TeV/y to PB/y, small to large data processing needs, integration of external resources (HPC), workflow synchronisation turnarounds (data quality monitoring)



A comparison of the yearly data volumes of current and future projects:



Credits: APS/[Alan Stonebraker](#) and V. Gülzow/DESY

Astroparticle Physics and Radio Astronomy

Knowledge transfer and expertise sharing among large ESFRIs on scientific computing

Data Recording from remote sources



Computing at CERN: the big picture

Data

Detector

Simulated Data

Interaction rate
~1 GHz

Bunch crossing rate
40 MHz

LEVEL-1 TRIGGER
<75 (100) kHz

Regions of Interest

LEVEL-2 TRIGGER
~3.5 kHz

EVENT FILTER
~200 Hz

Data Recording

Simulation (int. w/ det.)

Digitization (simul detector resp)

GEANT4

CPU **GPU** **HPCs**

updates > briefing > ATLAS event selection system readies for LHC Run 3

60 GeV

60 GeV

Computing at CERN - International Teachers Programme - xavier.espinal@cern.ch



From the Hit to the Bit: DAQ



CMS Experiment at the LHC, CERN

Data recorded: 2015-Sep-28 06:09:43.129280 GMT

Run / Event / LS: 257645 / 1610868539 / 1073



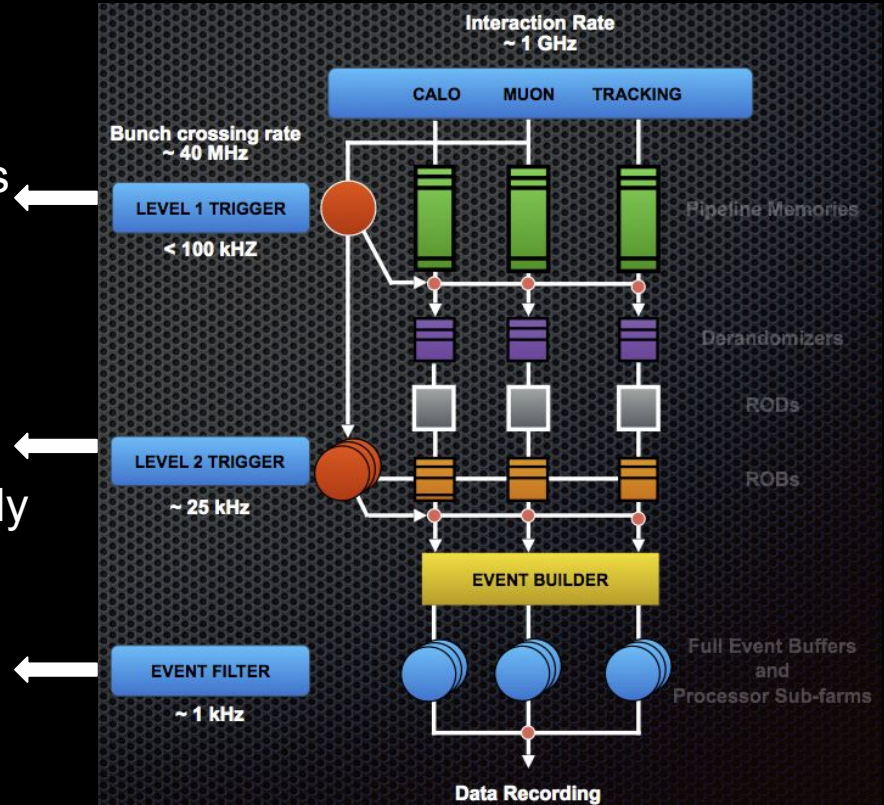
100 million channels

40 million pictures a second

Synchronised signals from all detector parts

From the Hit to the Bit: event filtering

- L1: 40 million events per second
 - Fast, simple information
 - Hardware trigger in a few microseconds
- L2: 100 thousand events per second
 - Fast algorithms in local computer farm
 - Software trigger in <1 second
- EF: Few 100 per second recorded for study




From the Hit to the Bit: event filtering

- L1: 40 million events per second

1PB/s: few hundred trillion euros/yr !!!
From 40 MHz to 100kHz

- L2: 100 thousand events per second
 - Fast algorithms in local computer farm
 - Software trigger in <1 second

- EF: Few 100 per second recorded for study



From the Hit to the Bit: event filtering


- L1: 40 million events per second

1PB/s: few hundred trillion euros/yr !!!
From 40 MHz to 100kHz

- L2: 100 thousand events per second

We keep ~1 event in a million
From 100 kHz to 25kHz

- EF: Few 100 per second recorded for study



From the Hit to the Bit: event filtering

- L1: 40 million events per second


1PB/s: few hundred trillion euros/yr !!!
From 40 MHz to 100kHz

- L2: 100 thousand events per second

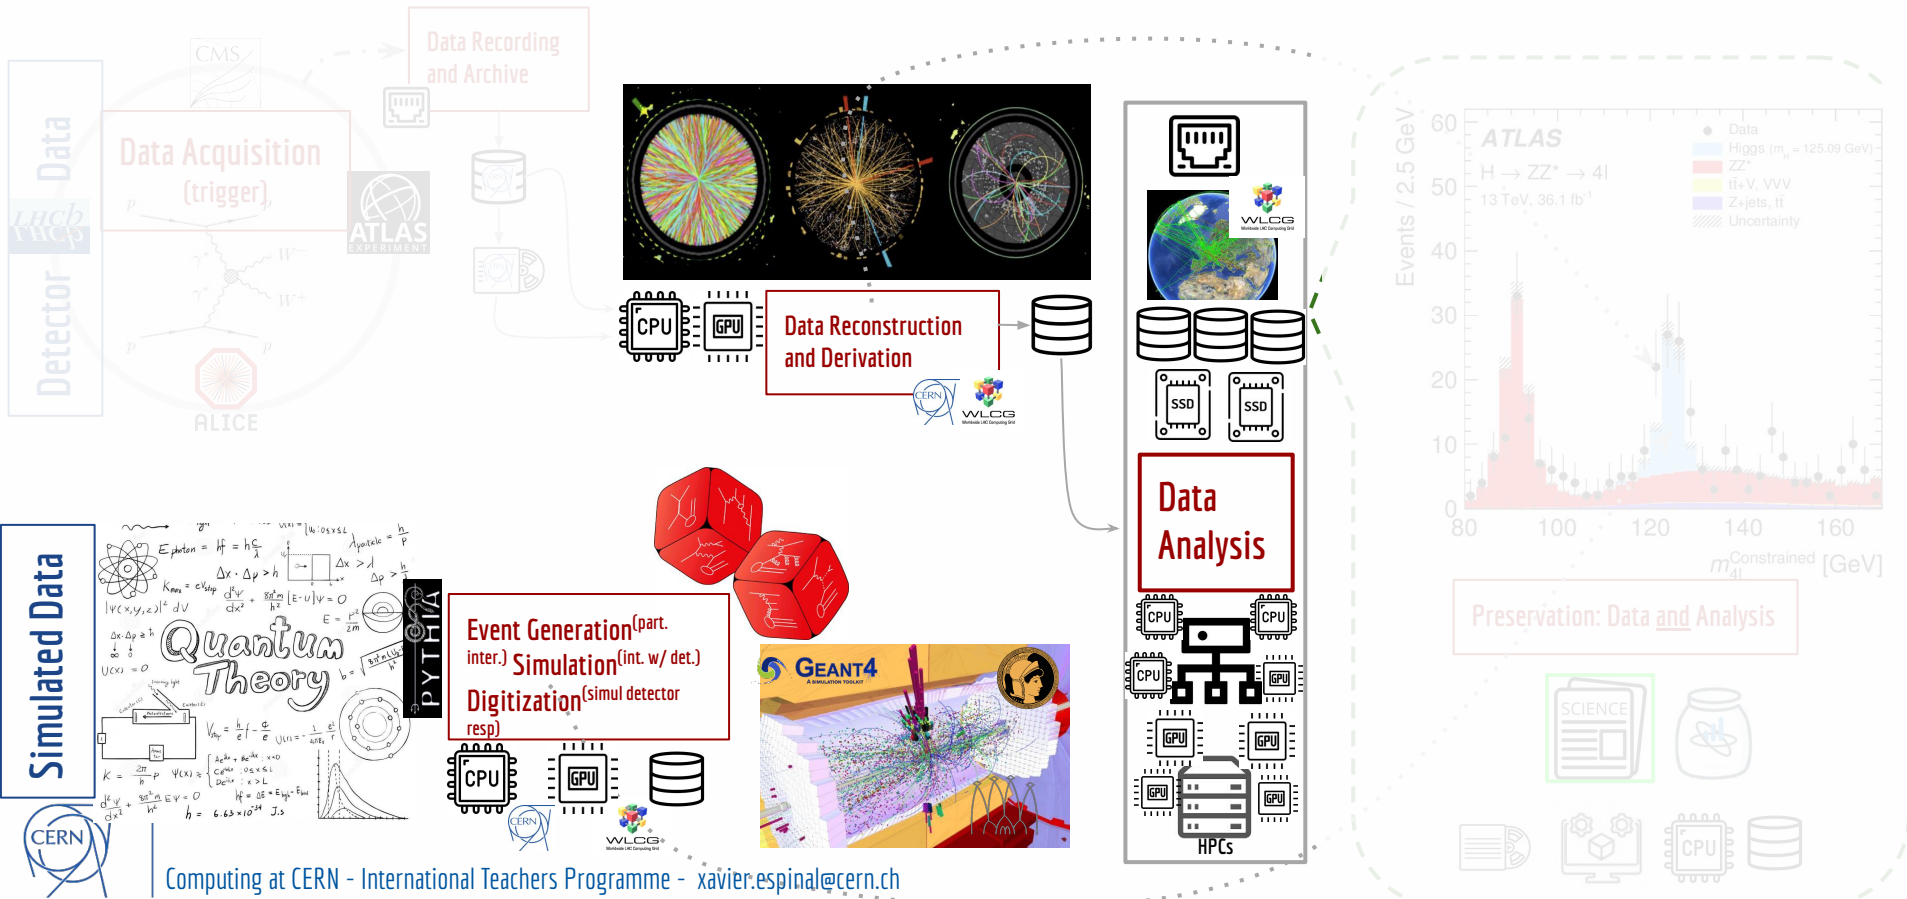
We keep ~1 event in a million
From 100 kHz to 25kHz

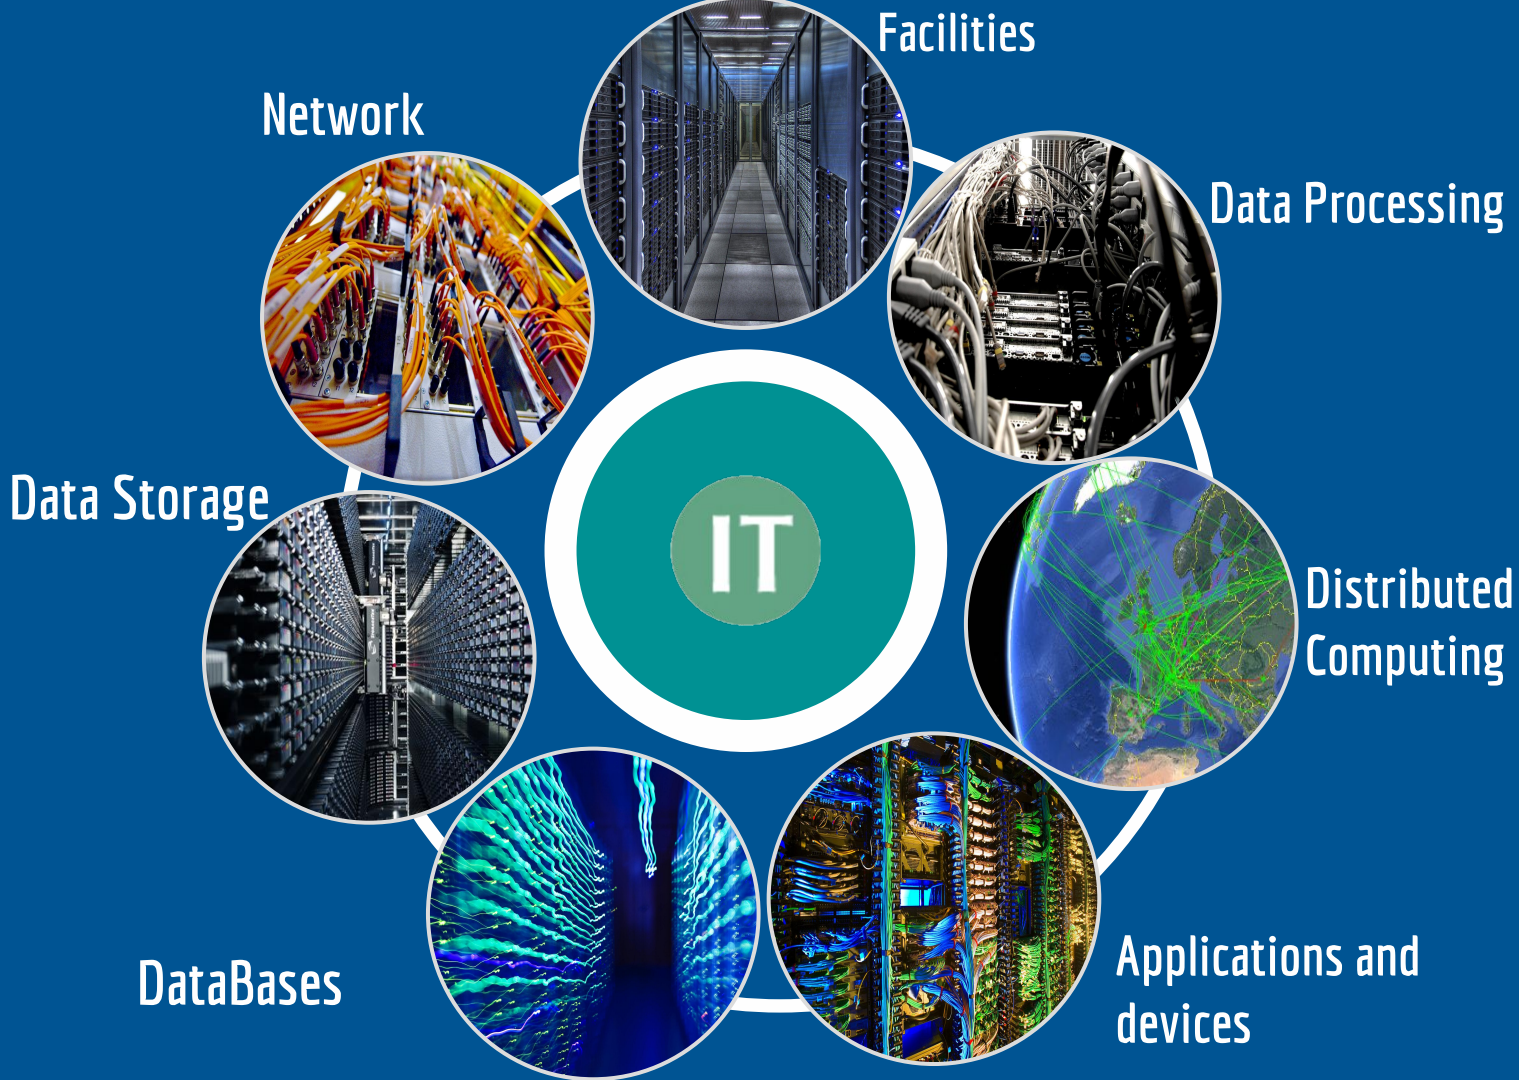
- EF: Few 100 per second recorded for study

From 25 kHz to few kHz
Final data rates ~10 GB/s





Computing at CERN: the big picture



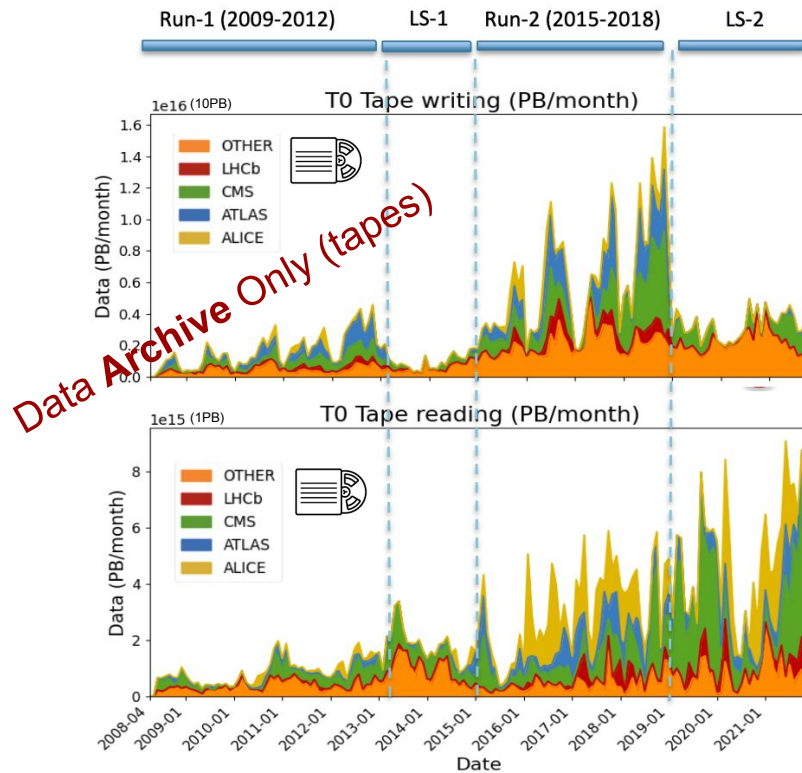




The Worldwide LHC Computing Grid



WLCG provides global computing resources to store, distribute and analyse the LHC data

The Worldwide LHC Computing Grid



The last LHC Run: data recording

- LHC Experiments recorded 88 Petabytes of data in 2018 (**15.8PB only in November**)
- The LHC data is aggregated at the CERN data centre to be stored, processed and distributed


Final update - Data - 2018



2018: 88 PB
ATLAS: 24.7
CMS: 43.6
LHCb: 7.3
ALICE: 12.4

inc. parked b-physics data

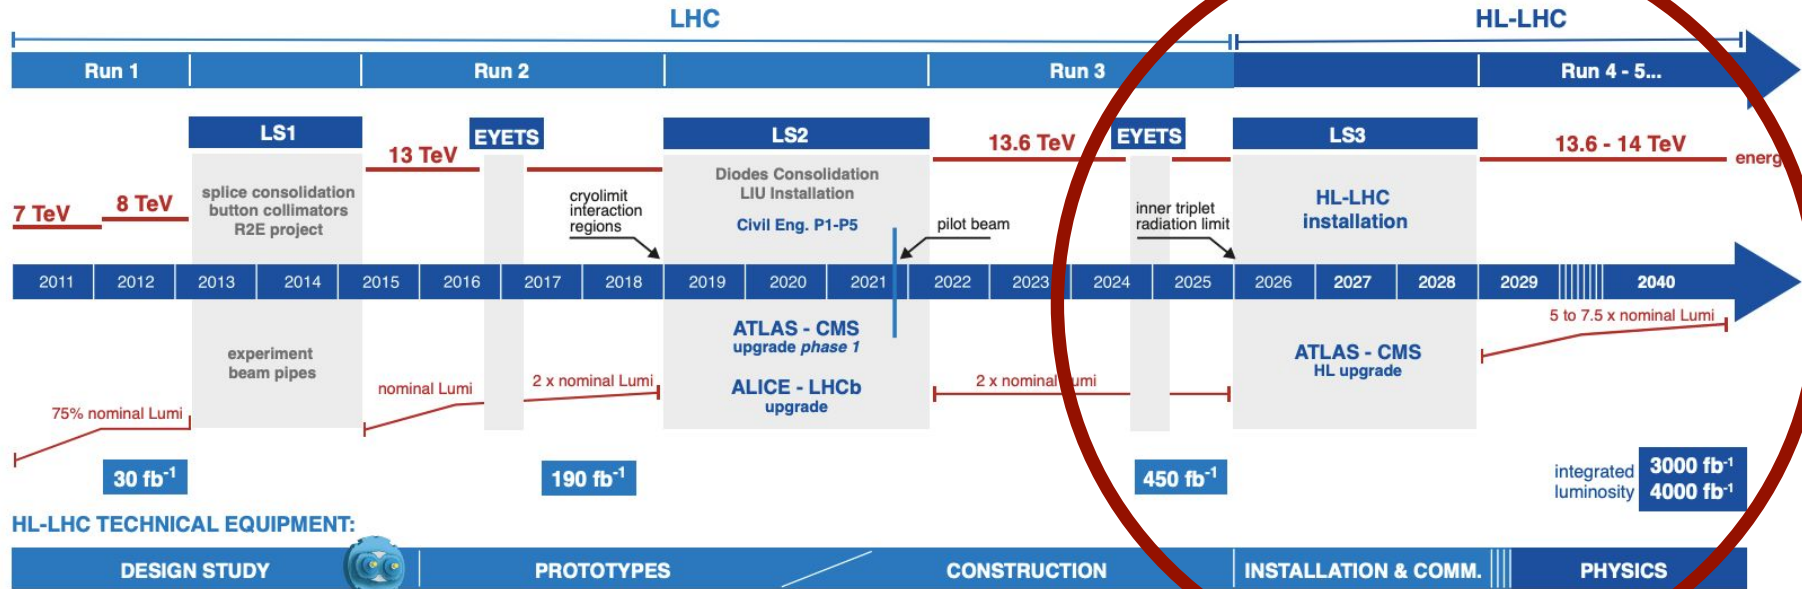
Ian Bird@WLCG-OB



HL-LHC: a computing challenge



LHC / HL-LHC Plan




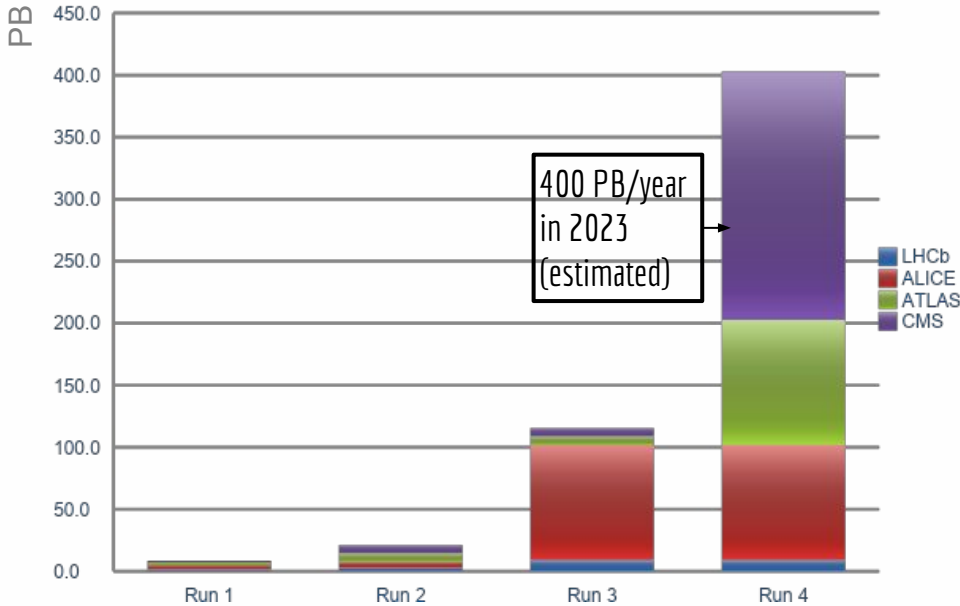
HL-LHC TECHNICAL EQUIPMENT:



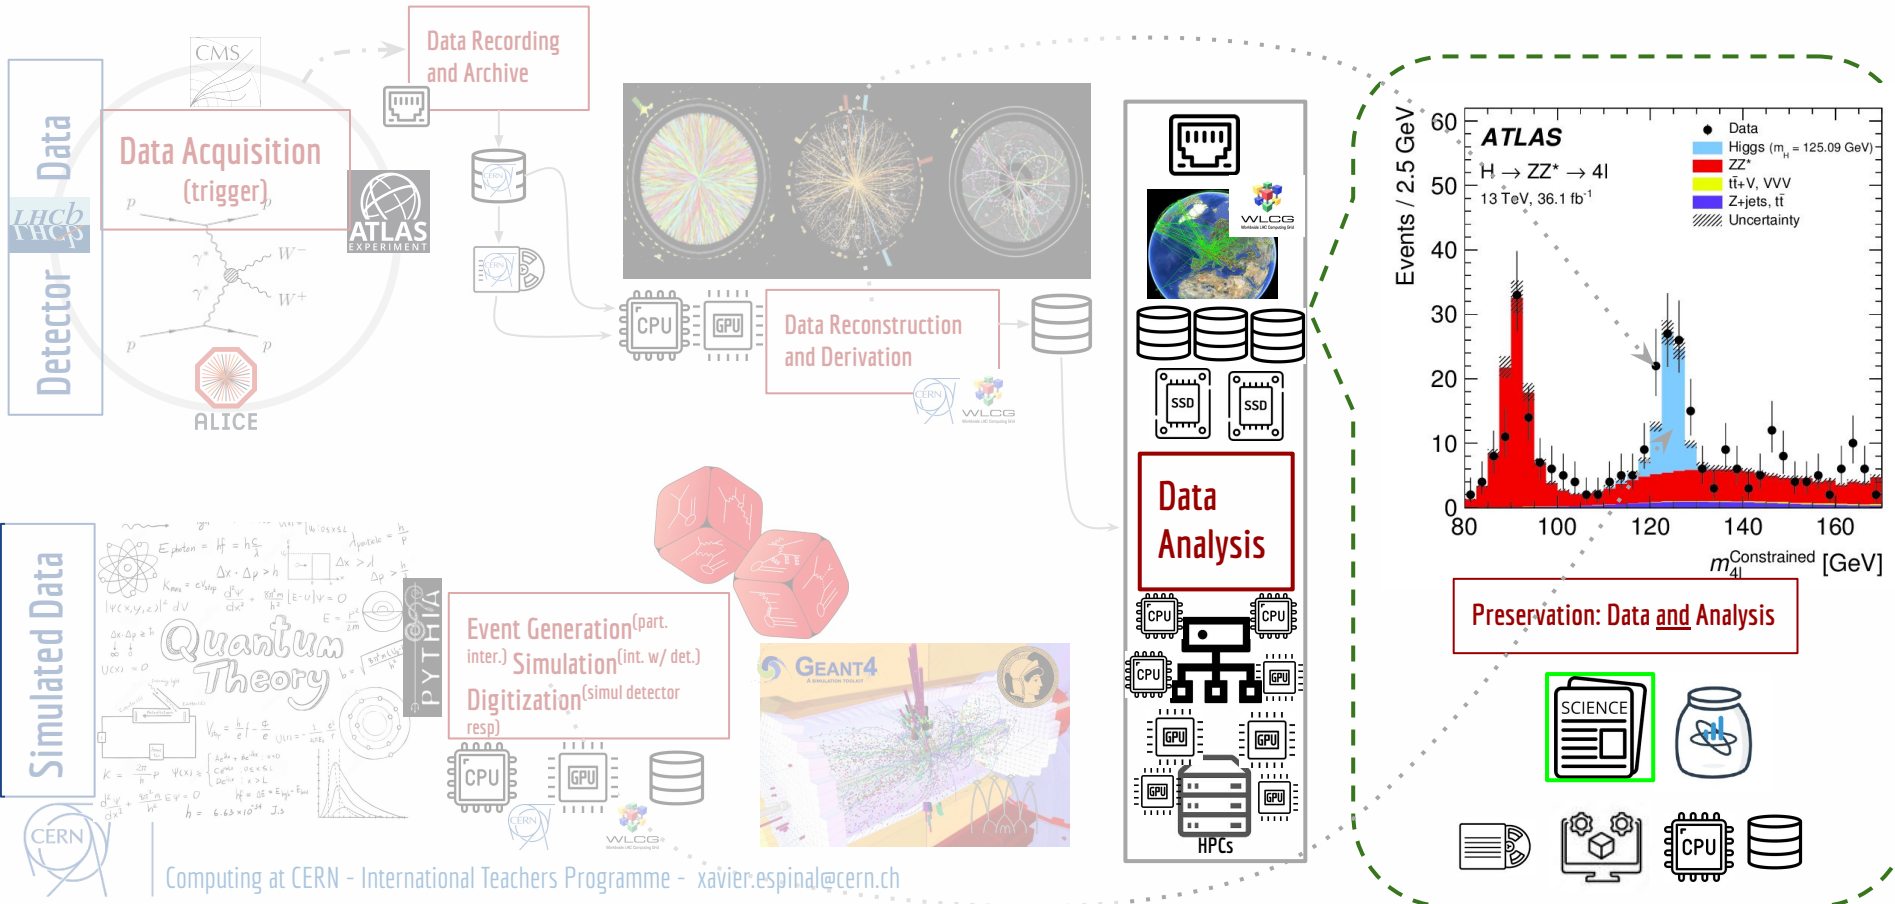
HL-LHC CIVIL ENGINEERING:



HL-LHC: a computing challenge




Computing at CERN: the big picture



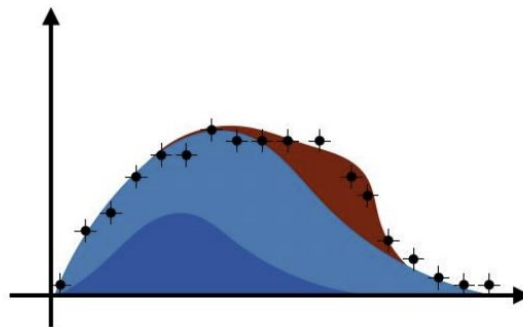
Data processing... and discovery!

Spawning a large cloud computing cluster at CERN for data processing using notebooks (70TB spread over 20k files). All processed in few minutes, real time!

Baseline only



Baseline +
New Physics



Baseline cannot describe data

... but baseline + new physics theory does -> Discovery!

Heinrich, Rocha @Kubecon and @CERN-ITTF

Keynote: Re-performing a Nobel Prize Discovery on Kubernetes

<https://www.youtube.com/watch?v=CTfp2woVEkA> (10:15 start, 14:30 populating plot)

Da


very!

Spawning a
over 20k file

Baseline on



ooks (70TB spread



Hein
Key

<https://www.youtube.com/watch?v=CTfp2woVEkA> (10:15 start, 14:30 populating plot)




CERN Data Center


- Built in the 70s on the CERN site (Meyrin, Geneva)
 - 3.5 MW for equipment
- Hardware generally based on commodity
 - 10,000 servers, providing 500,000 processor cores (14.000 VMs)
 - 120,000 disk drives providing 0.6EB disk space (1EB=1000PB)
 - 40,000 tapes drives, providing 0.5EB capacity (1EB=1000PB)



CERN Data Center



Green IT



Power Usage Effectiveness (PUE): Total Facility Energy / IT Facility Energy



Building work for CERN's new data centre in Prévessin begins

On Friday, 22 April, a ceremony was held to mark the beginning of construction of CERN's new energy-efficient data centre

22 APRIL, 2022 | By Andrew Purcell



Representatives of CERN and of the EQUANS France-Léon Grosse-Agapé consortium involved in the first-stone ceremony



An artist's impression of the CERN Data Centre in Prévessin. The facility will consist of three floors, to be progressively filled with computing equipment over the first ten years of operation. (Credit: + IMGs - Rocco Valentines)

Operational in the second half of 2023

“Computing is central to CERN’s mission,” says Charlotte Warakaulle, CERN Director for International Relations, who participated in the first-stone ceremony for the new data centre. “It turns data into knowledge, helping physicists unlock the secrets of the universe.”

The CERN Data Centre in Prévessin will provide computing resources up to a total electrical power requirement of 12 megawatts. These resources will be delivered in three phases. Each phase corresponds to one of the three floors of the new data centre, with the first phase set to run from 2023 to 2025. It will see computing resources requiring up to 4 megawatts of electrical power installed; this is approximately the same as the power of the current CERN Data Centre in Meyrin for computing (excluding cooling).

Take-away (1/3)

- Computing is instrumental for science. Detectors and sensors evolving: growing IT demands.
- LHC raw data rates are PB/s scale but lowered to GB/s after data *filtering*
- 90PB of LHC data in 2018 (15.8PB in Nov only)
- 1EB data transferred world-wide
- Scientific data already at the **Exabyte scale**:
 - 1EB = 1.000PB = 1.000.000 TB = 1.000.000.000 GB
(1TB is your computer) (100 GB is your smartphone)

Take-away (2/3)

- Power and Heat management: PUE and Green-IT
- Data centers run on commodity hardware
- Big computing companies dominating the market: G, MS, DB, FB,...
- CERN remains largest scientific repository in the world



Take-away (3/3)

- High-Luminosity LHC brings new challenges in computing: time for **new ideas** and **R&D!**
- **Fundamental** science continue to be the main inspiration for **revolutionary** ideas, due to revolutionary needs
- Industry has well defined offer and demand. We do not. This is the key for **innovation**
- ...and **innovation** foster technological advancements that percolates to the society




PROTON PHYSICS: STABLE BEAMS

Energy: **6800 GeV** **I B1:** **2.96e+11** **I B2:** **2.96e+11**


Beta* IP1: **0.30 m** Beta* IP2: **10.00 m** Beta* IP5: **0.30 m** Beta* IP8: **2.00 m**

Inst. Lumi [(ub.s)⁻¹] IP1: **9.57** IP2: **0.55** IP5: **10.65** IP8: **2.93**

FBCT Intensity and Beam Energy Updated: 10:37:41



Instantaneous Luminosity Updated: 10:37:43



Thank you!

Comments (06-Jul-2022 05:09:17)

**** STABLE BEAMS @ 6.8 TeV ****

XRPs IN
beta* levelling ON

AFS: Single_3b_2_2_2

BIS status and SMP flags

Link Status of Beam Permits

B1

B2

true **true**

Global Beam Permit

true **true**

Setup Beam

false **false**

Beam Presence

true **true**

Moveable Devices Allowed In

true **true**

Stable Beams

true **true**

PM Status B1 **ENABLED**

PM Status B2 **ENABLED**