Phase VI Challenges and Opportunities

Alberto Di Meglio – CERN openlab Head

Monday, 24 June 2019

CERN: A UNIQUE ENVIRONMENT

Pushing technologies to their limits

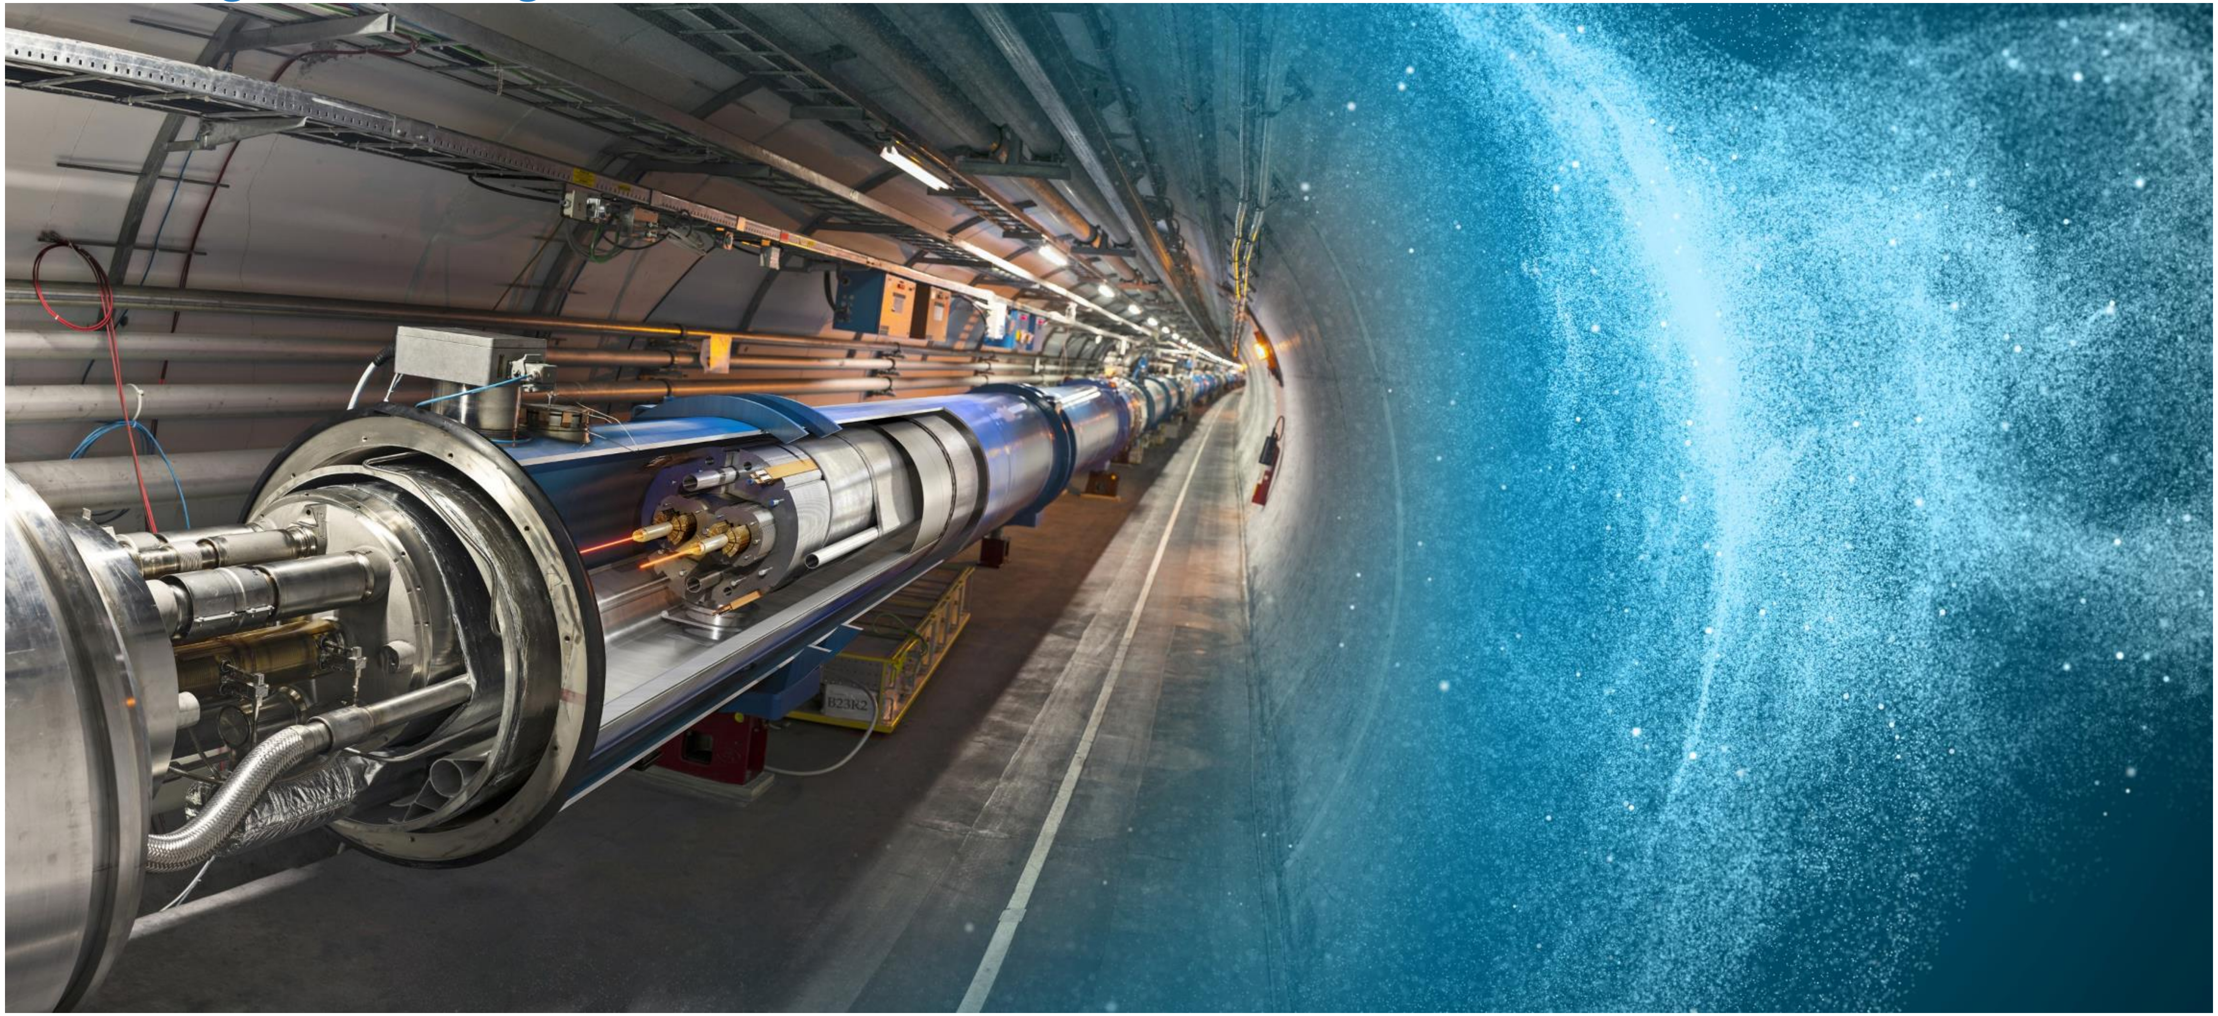

CERN

“Science for peace”

- International organisation close to Geneva, straddling Swiss-French border, founded 1954
- Facilities for fundamental research in particle physics
- 23 member states, 1.1 B CHF budget
- ~ 3’500 staff, fellows, apprentices, …
- ~ 14’400 associates

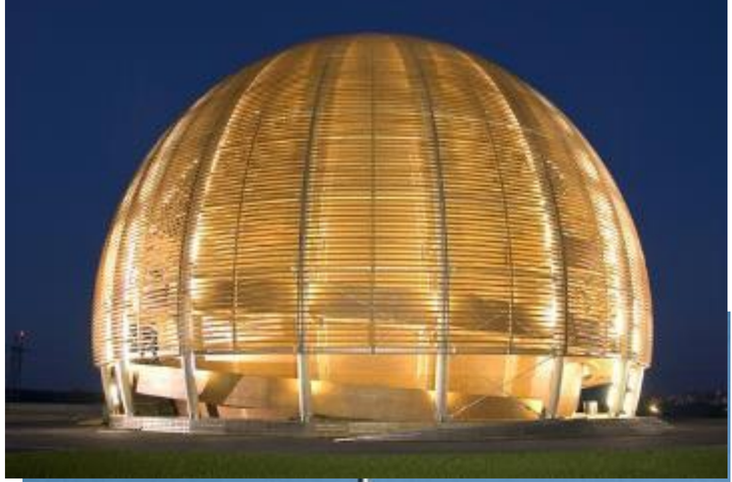

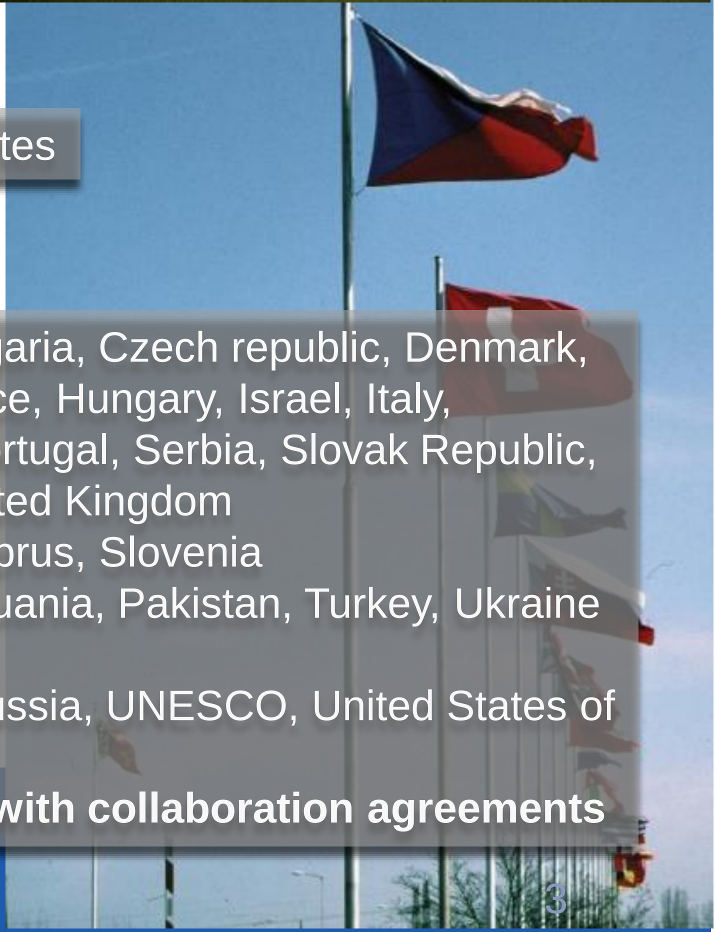

1954: 12 Member States

Members: Austria, Belgium, Bulgaria, Czech republic, Denmark, Finland, France, Germany, Greece, Hungary, Israel, Italy, Netherlands, Norway, Poland, Portugal, Serbia, Slovak Republic, Spain, Sweden, Switzerland, United Kingdom
Candidate for membership: Cyprus, Slovenia
Associate members: India, Lithuania, Pakistan, Turkey, Ukraine (Croatia)
Observers: EC, Japan, JINR, Russia, UNESCO, United States of America
Numerous **non-member states with collaboration agreements**


CERN
CMS
CMS
1 PB/sec
> 2000 disks/sec
LHCb
LHCb
ALICE
ALICE
ATLAS
ATLAS

Storage, Reconstruction, Simulation, Distribution

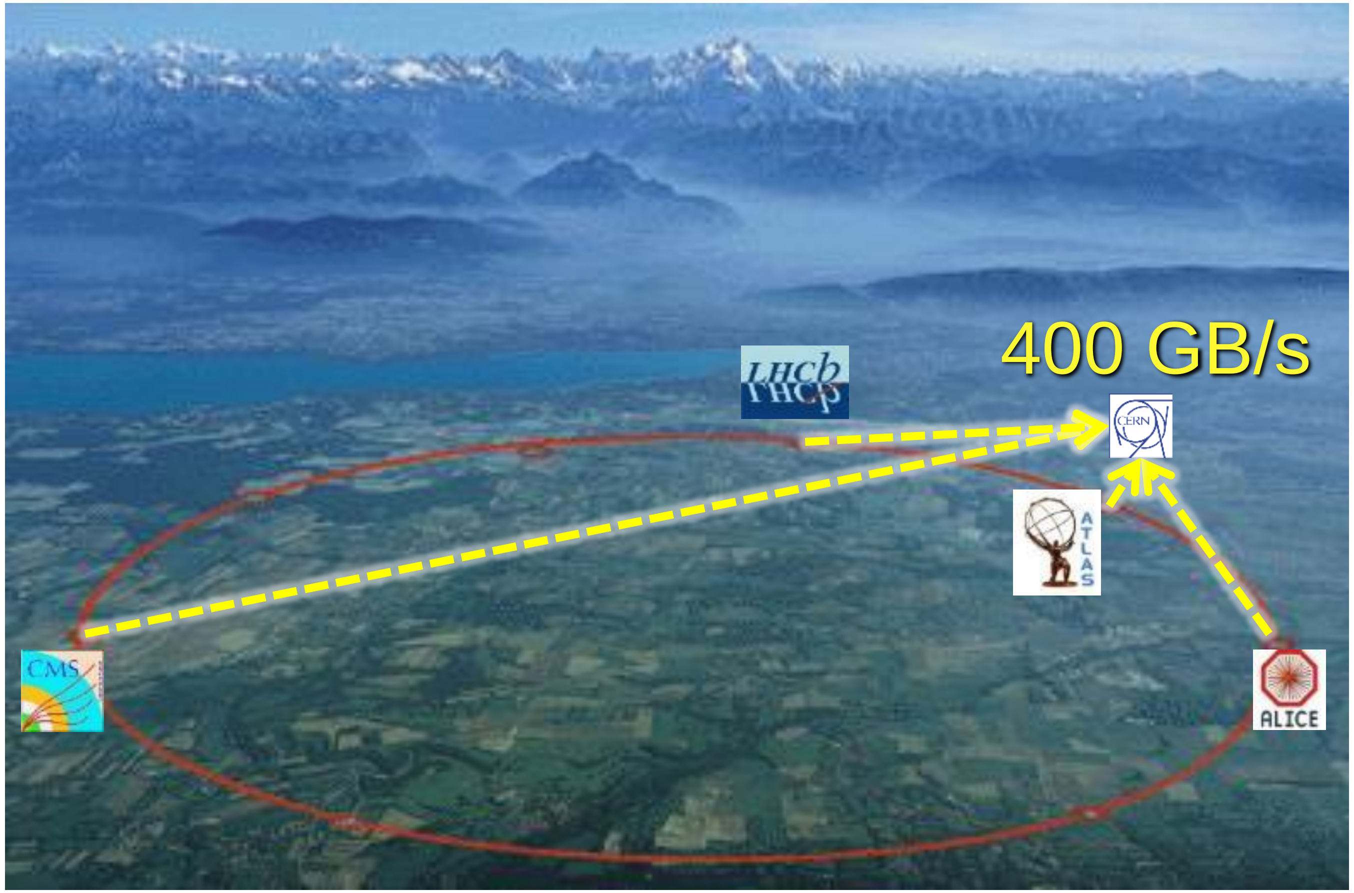

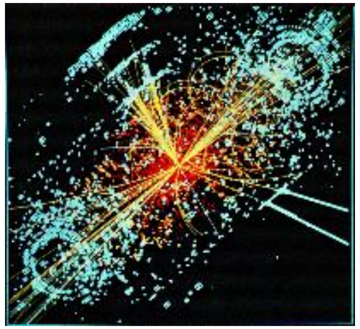

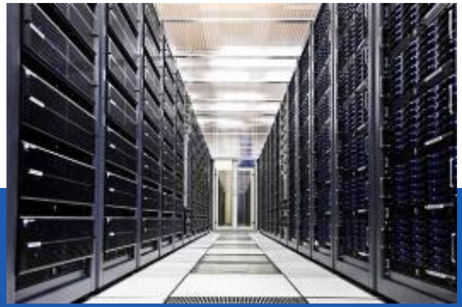

The CERN Data Centre in Numbers

15 000 Servers

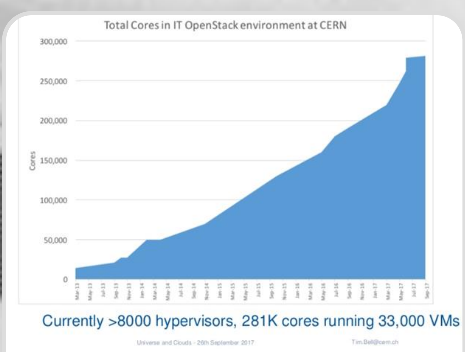

280 000 Cores

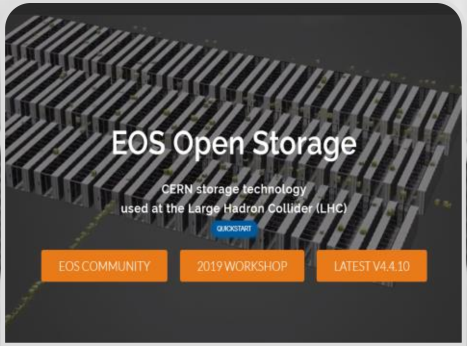

280 PB Hot Storage

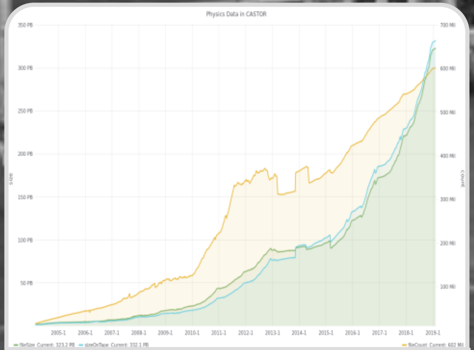

350 PB Cold Storage

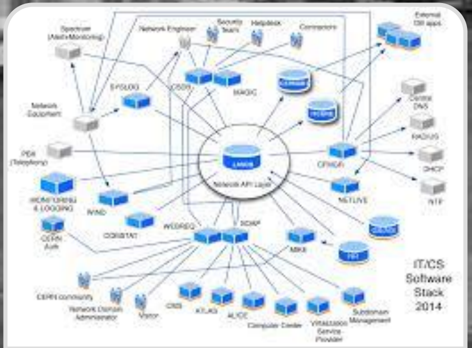

35 000 km Fiber Optics

Worldwide LHC Computing Grid

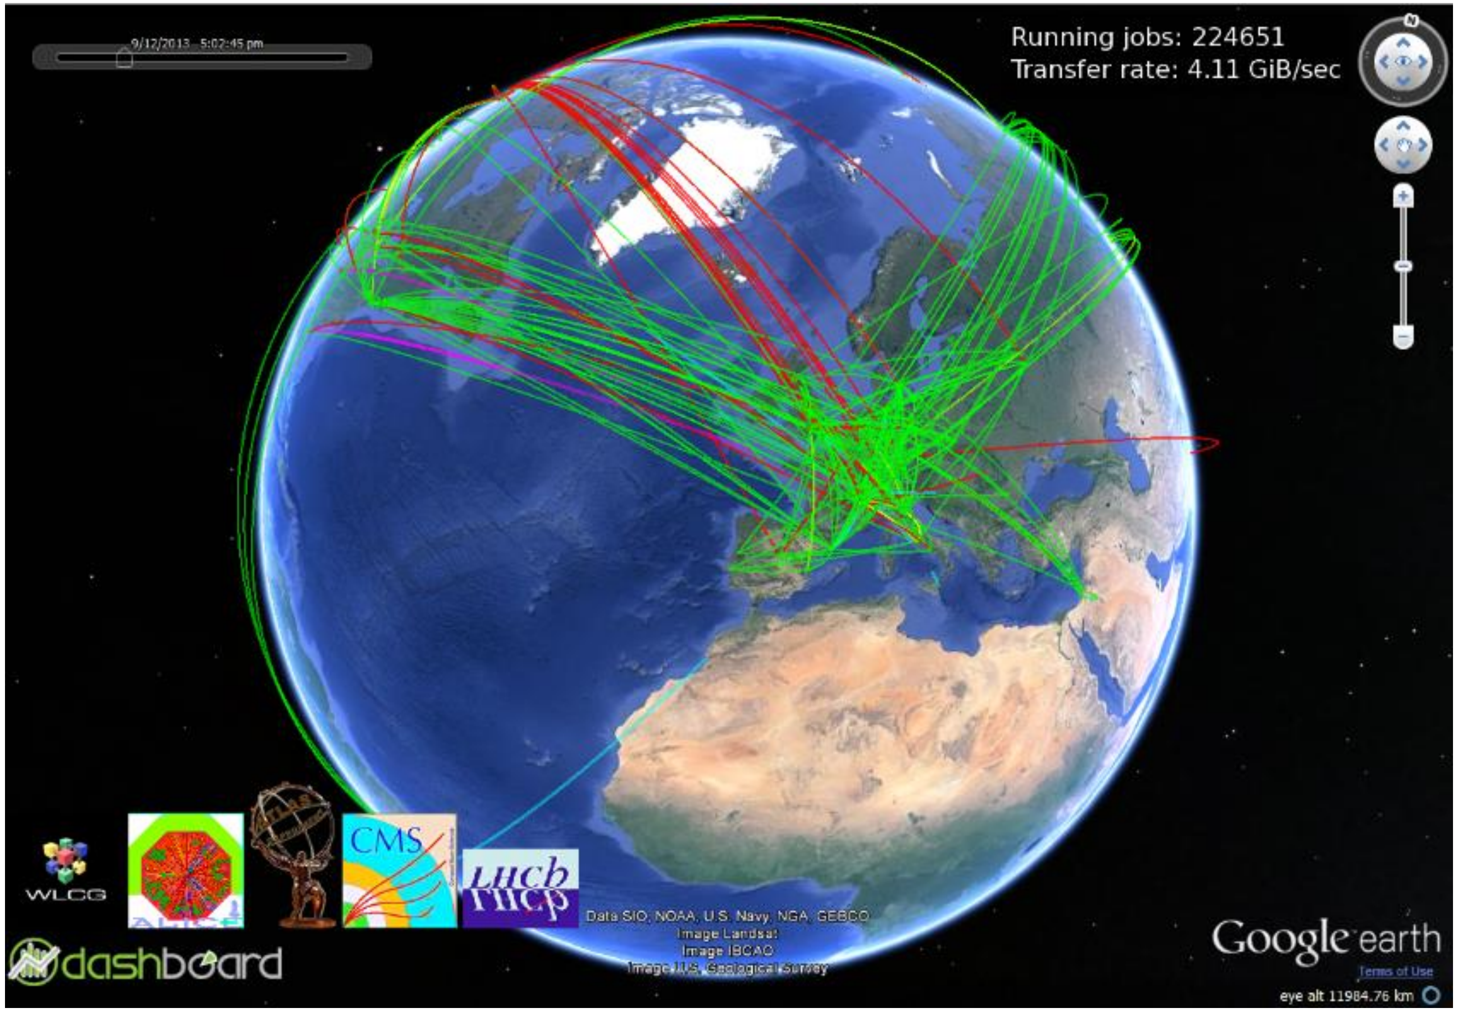

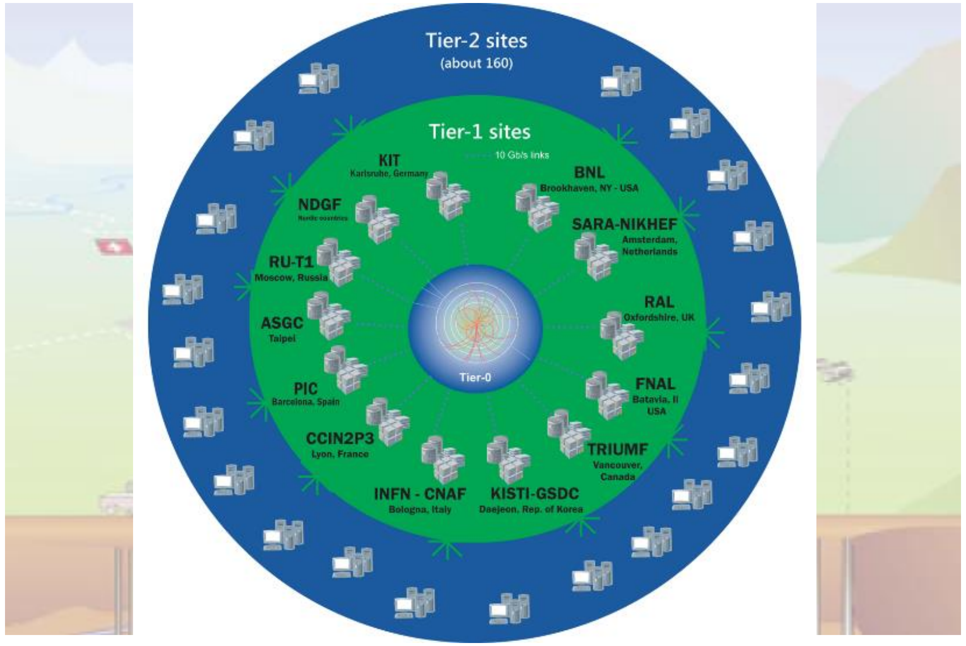

Tier-0 (CERN):
- Data recording
- Initial data reconstruction
- Data distribution

Tier-1 (14 centres):
- Permanent storage
- Re-processing
- Analysis

Tier-2 (72 Federations, ~149 centres):
- Simulation
- End-user analysis

- 760,000 cores
- 700 PB

The Higgs Boson

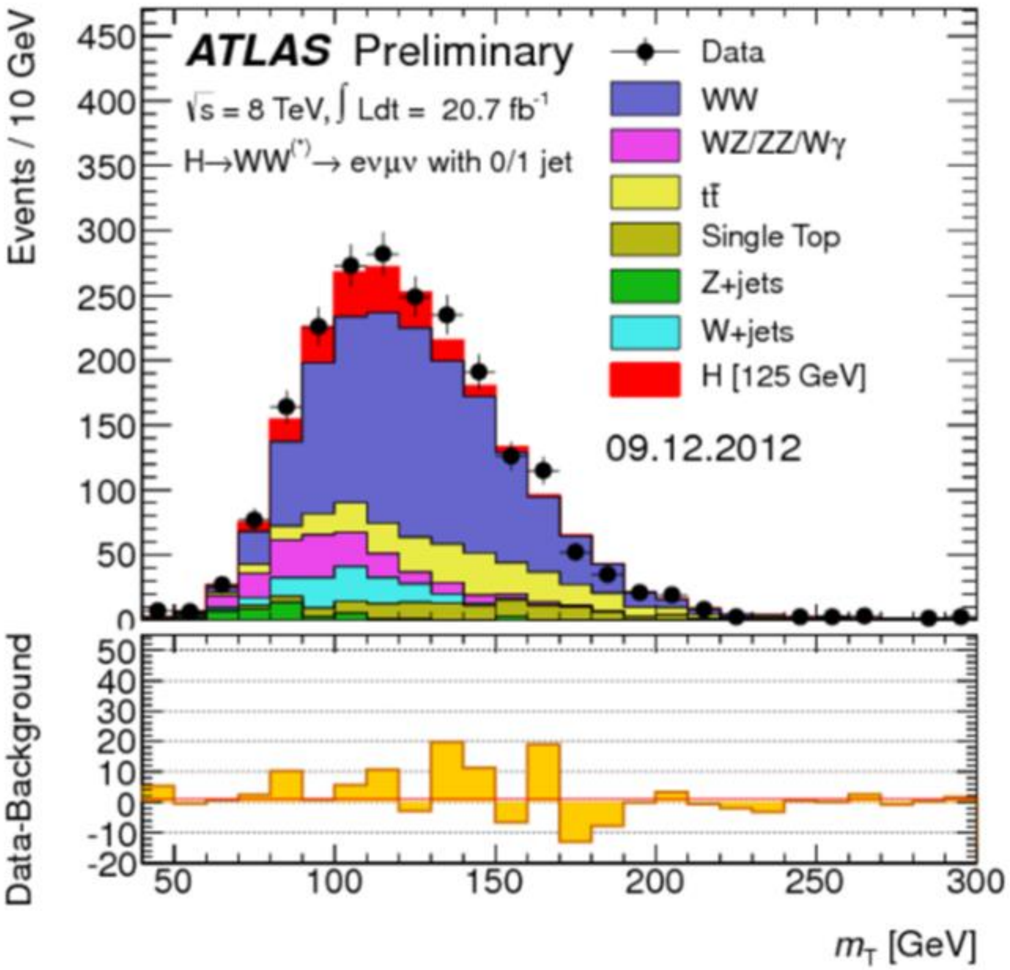

The Higgs Boson completes the Standard Model,
but the Model explains only about 5% of our Universe

What is the other 95% of the Universe made of?

How does gravity really works?

Why there is no antimatter in nature?

LHC Schedule

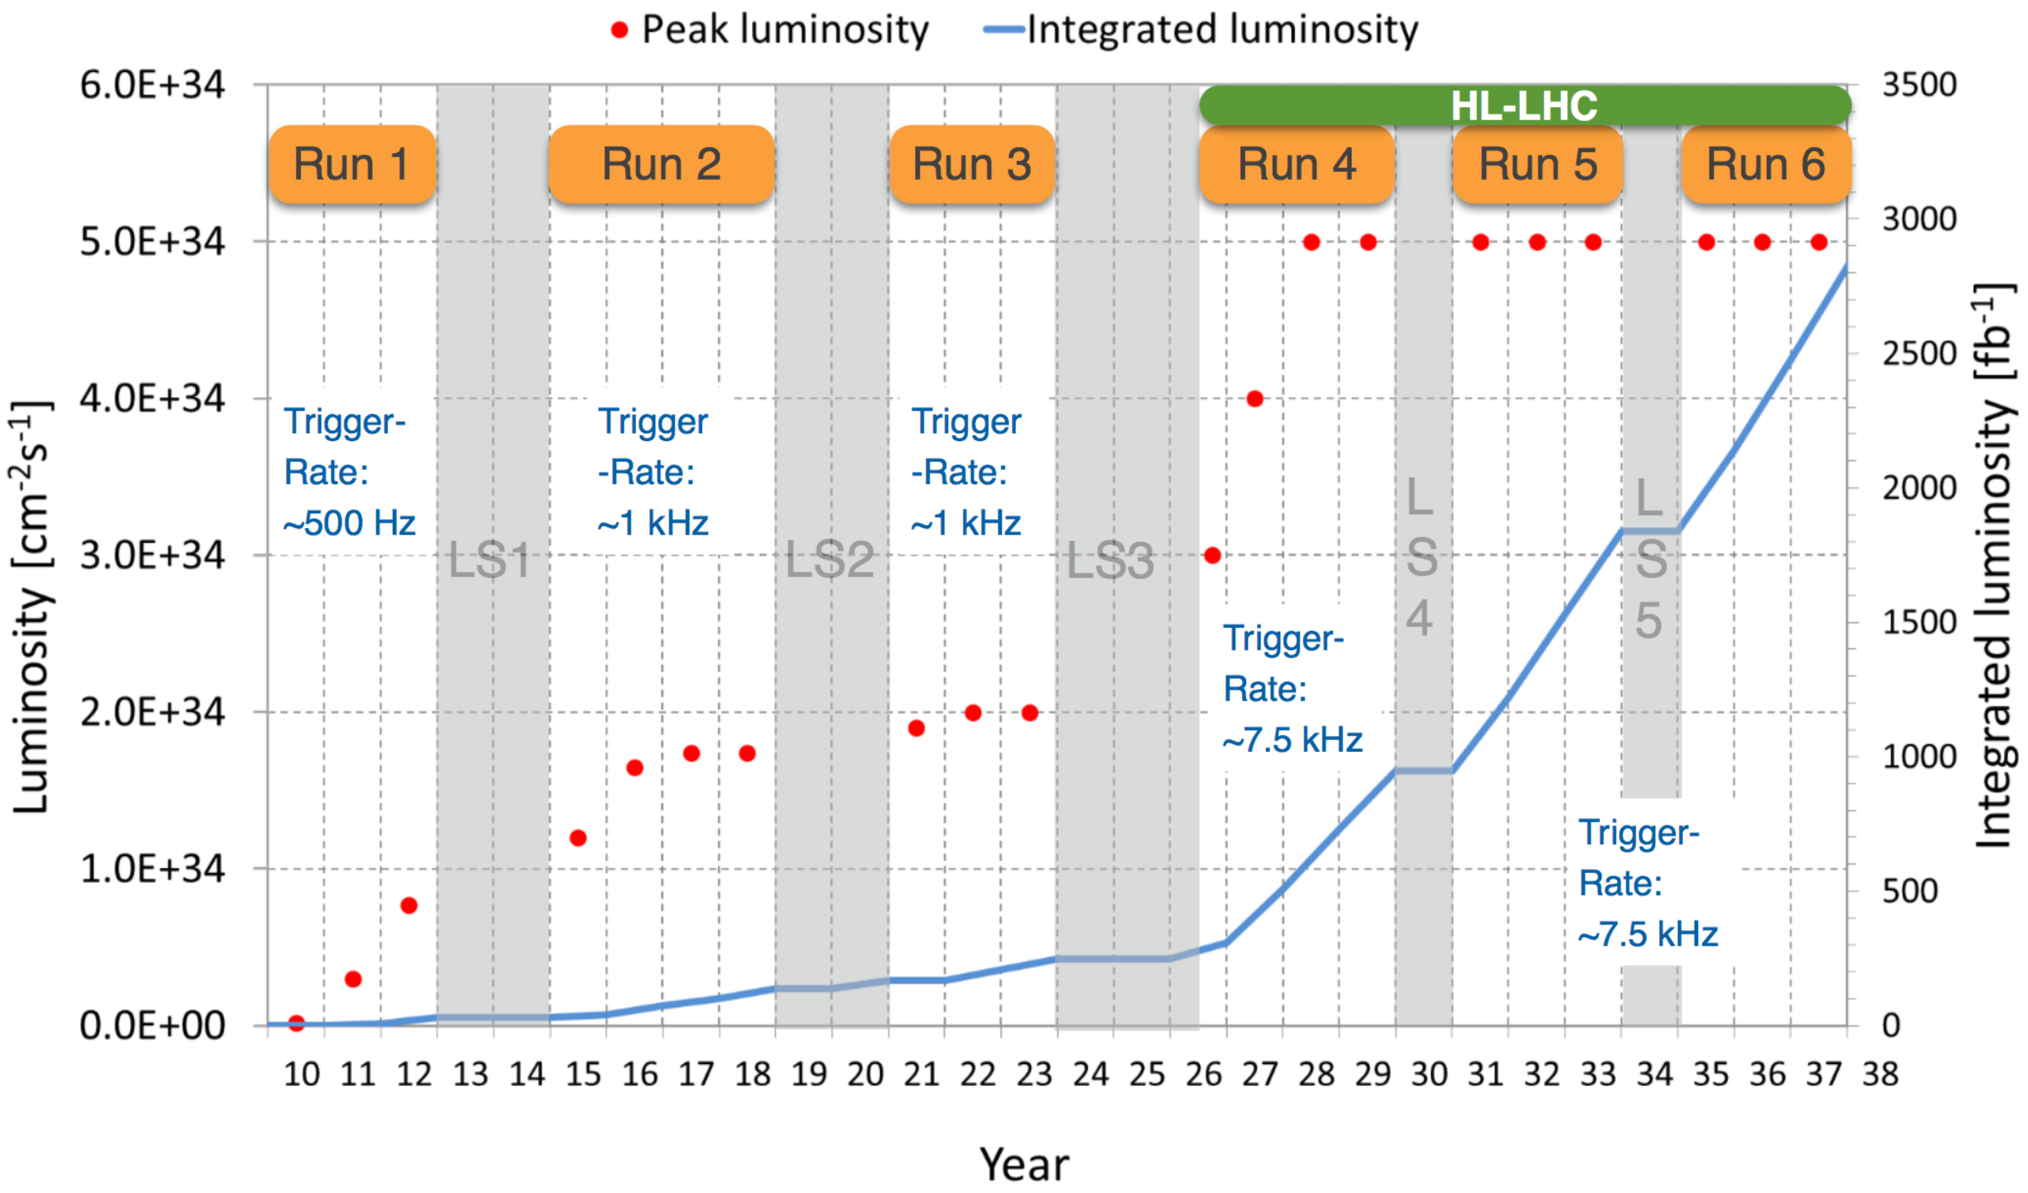

LHC Run3 and Run4

Scale and Challenges

2009 2010 2011 2012 2013 2014 2015 2016 2017 2018 2019 2020 2021 2022 2023 2024 2025 … … 2030?

First run | LS1 | Second run | LS2 | Third run | LS3 | HL-LHC | FCC?

Raw data volume for LHC increases exponentially and with it processing and analysis load

Technology at ~20%/year will bring x6-10 in 10-11 years

Estimates of resource needs at HL-LHC x10 above what is realistic to expect from technology with reasonably constant cost

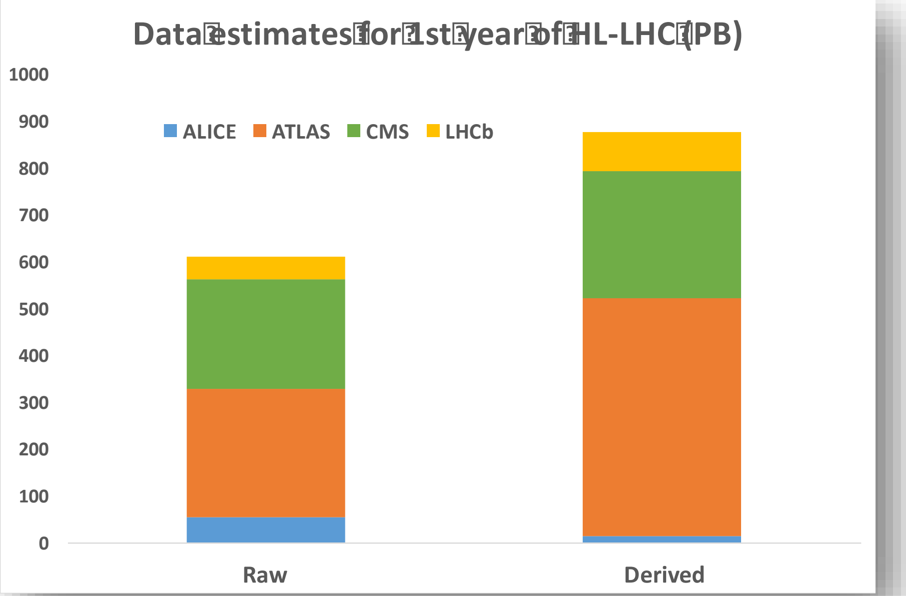

Data:
- Raw 2016: 50 PB → 2027: 600 PB
- Derived (1 copy): 2016: 80 PB → 2027: 900 PB

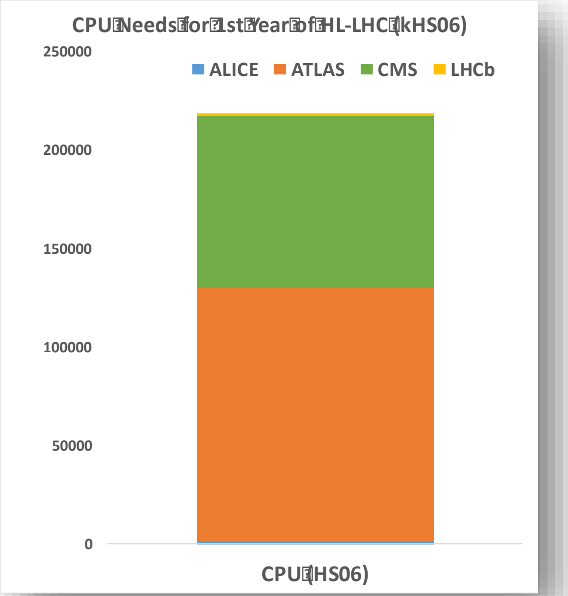

CPU:
- x60 from 2016

Technology revolutions are needed

CERN openlab

DRIVING INNOVATION SINCE 2001

We are now in our sixth three-year phase

SET UP
2001

I
2003

II
2006

III
2009

IV
2012

V
2015

2018
Sixth Phase
VI

CERN openlab Collaboration Board 2018

CERN openlab Governance

The Collaboration Board is composed of the representatives of the member Companies and Research Institutes

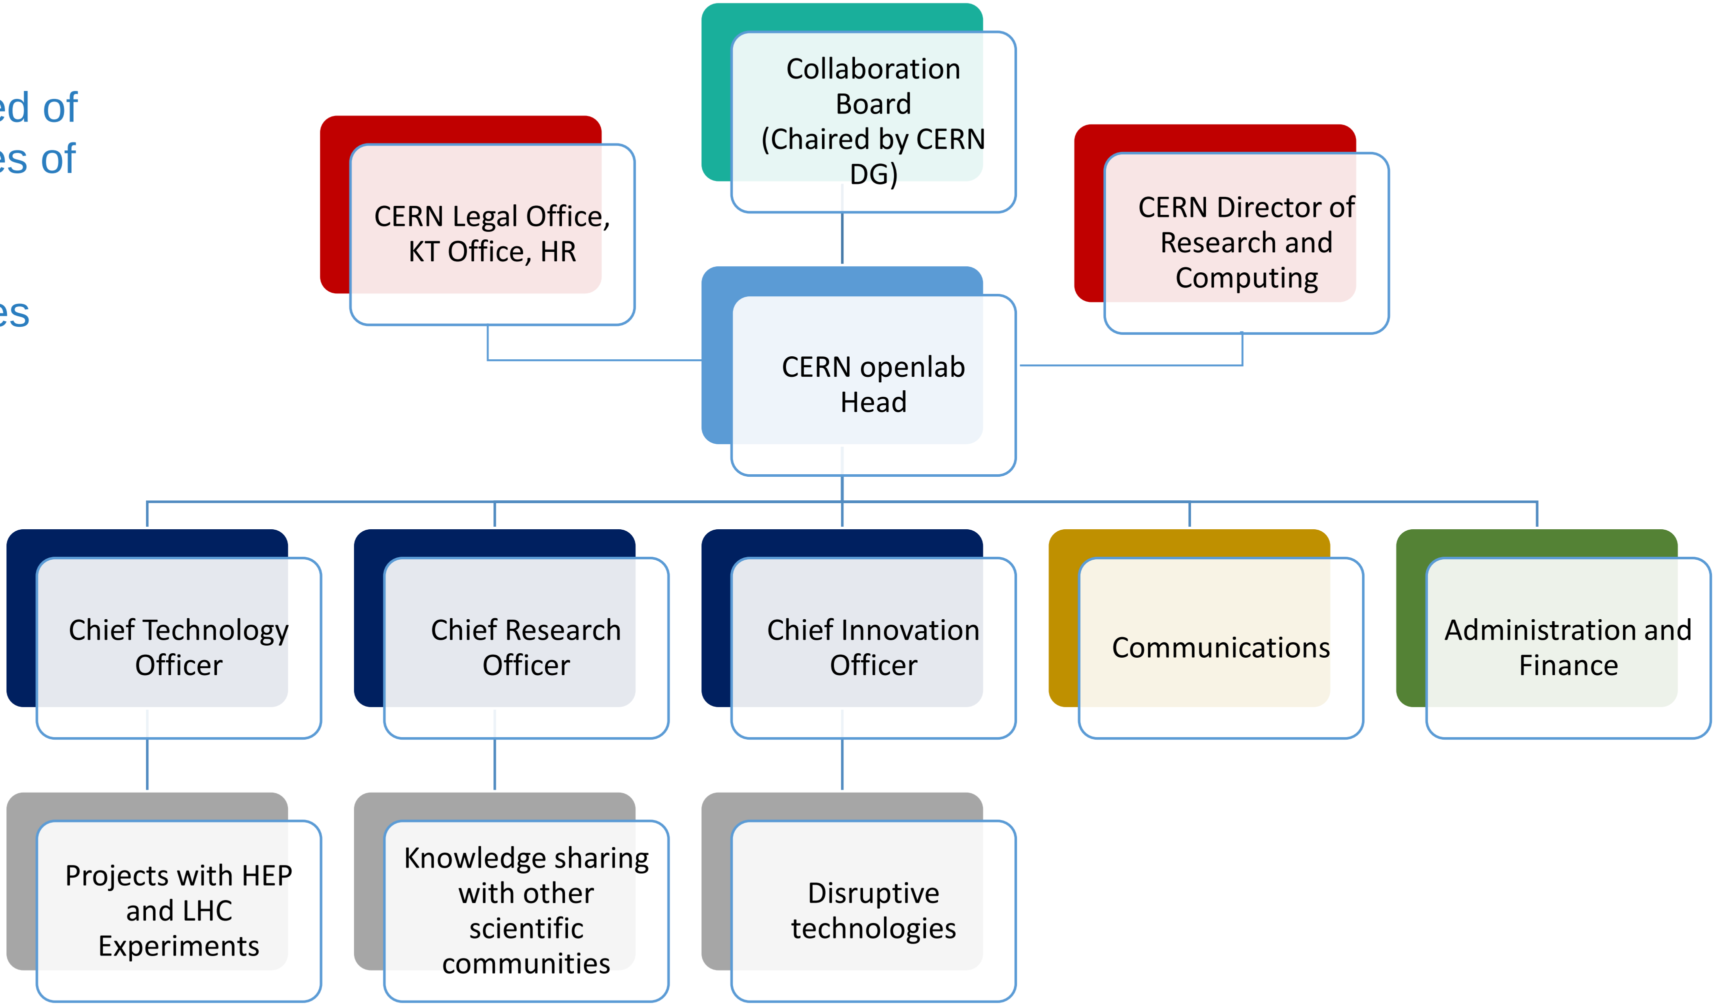

CERN OPENLAB'S MISSION

Our recipe for success

Evaluate and test state-of-the-art technologies in a challenging environment and improve them in collaboration with industry.

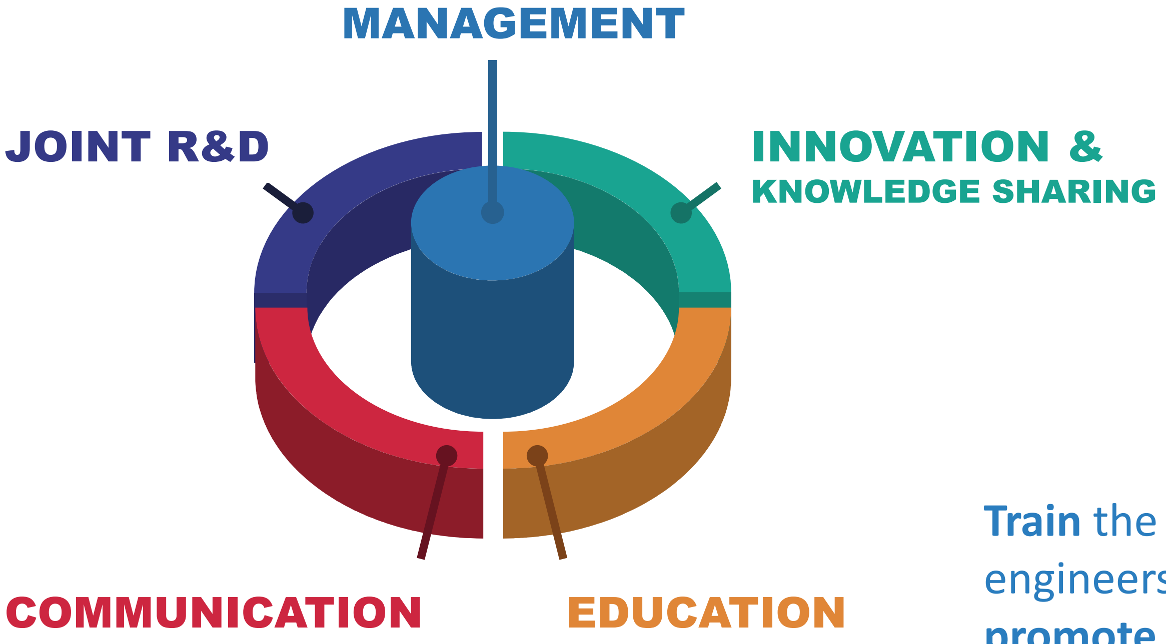

Collaborate and exchange ideas with other communities to create knowledge and innovation.

Communicate results, demostrate impact, and reach new audiences.

Train the next generation of engineers/researchers, **promote** education and cultural exchanges.

JOINT R&D PROJECTS (PHASE VI)

High-bandwidth fabrics, accelerated platforms for data acquisition

Fast simulation, Data quality monitoring, anomaly detection, physics data reduction, benchmarking/scalability, systems biology and large-scale multi-disciplinary platforms

Simulation, HPC on the Cloud, benchmarking

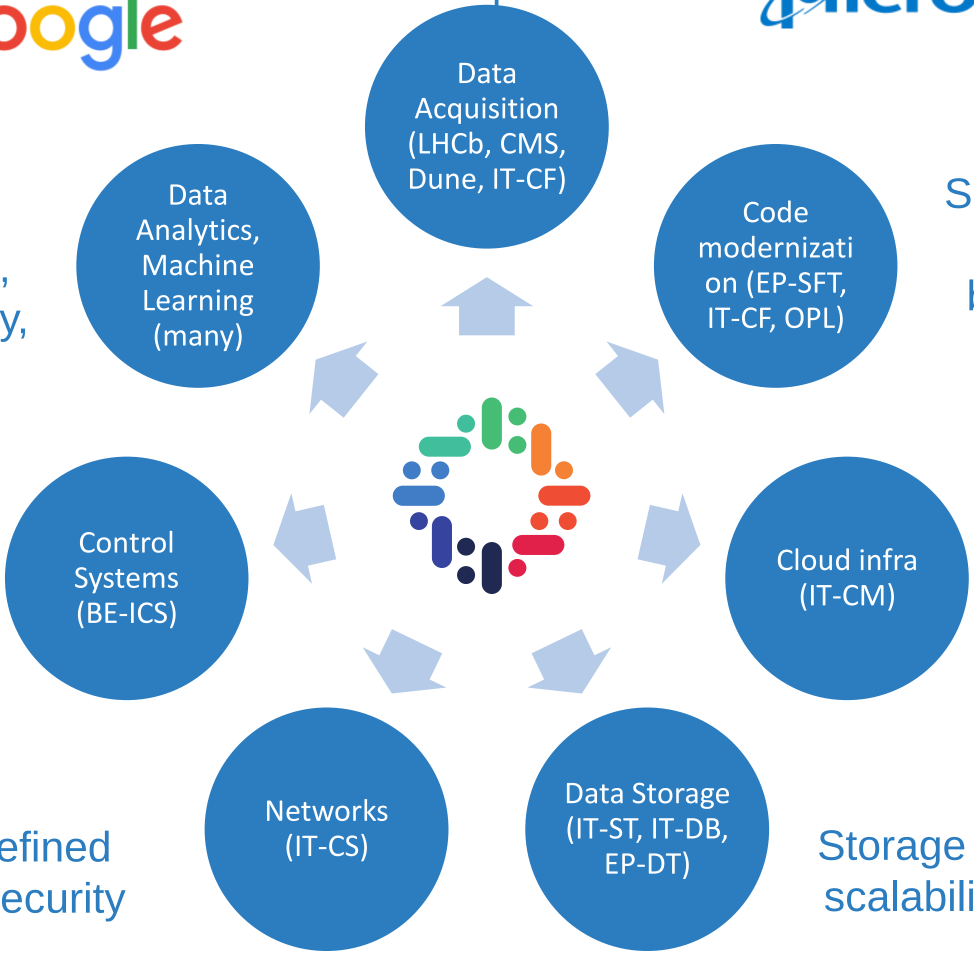

Predictive/proactive maintenance and operations

Cloud federations, containers, scalability

Software Defined Networks, Security

Storage architectures, scalability, monitoring

ICT WHITE PAPER 2017

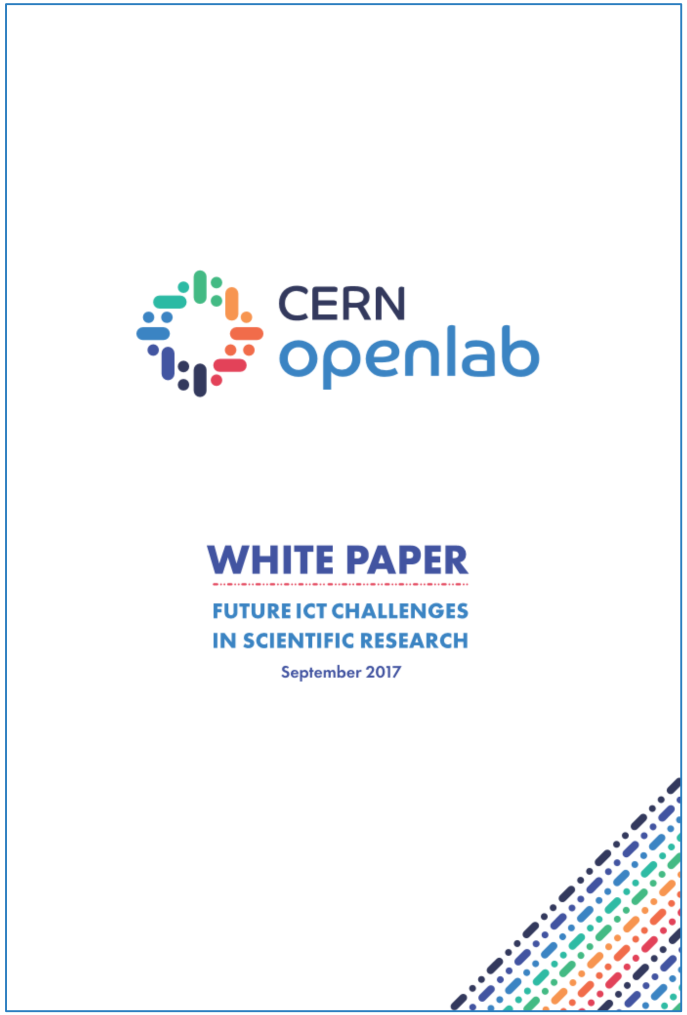
CERN openlab

WHITE PAPER

FUTURE ICT CHALLENGES IN SCIENTIFIC RESEARCH

September 2017

Published on 21 September 2017

Summarises the results of extensive consultations with CERN experts, LHC Experiments representatives, ICT companies and international research labs and academic institutes

Based on 4 major research topics and 16 challenge areas

A solid starting point for a very challenging and constructive Phase VI

Three Main Areas of R&D

High Luminosity LHC

COMPUTING CHALLENGES

Increase **data centre performance** with hardware accelerators (FPGAs, GPUs, ..) optimized software

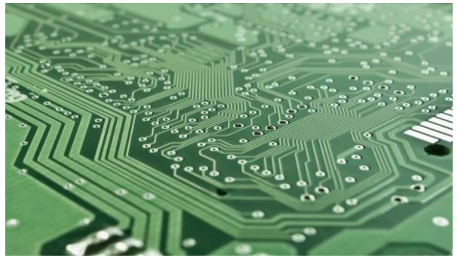

Scale out capacity with public clouds, HPC, new architectures

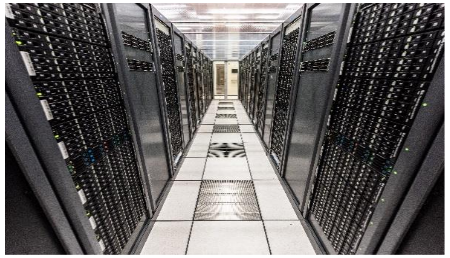

Change the computing paradigms with new technologies like **Machine Learning, Deep Learning**, Advanced Data Analytics, **Quantum Computing**

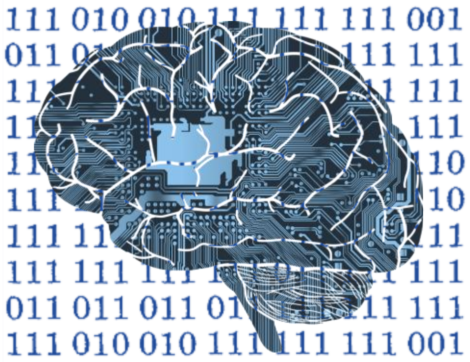

Computing and Data Challenges

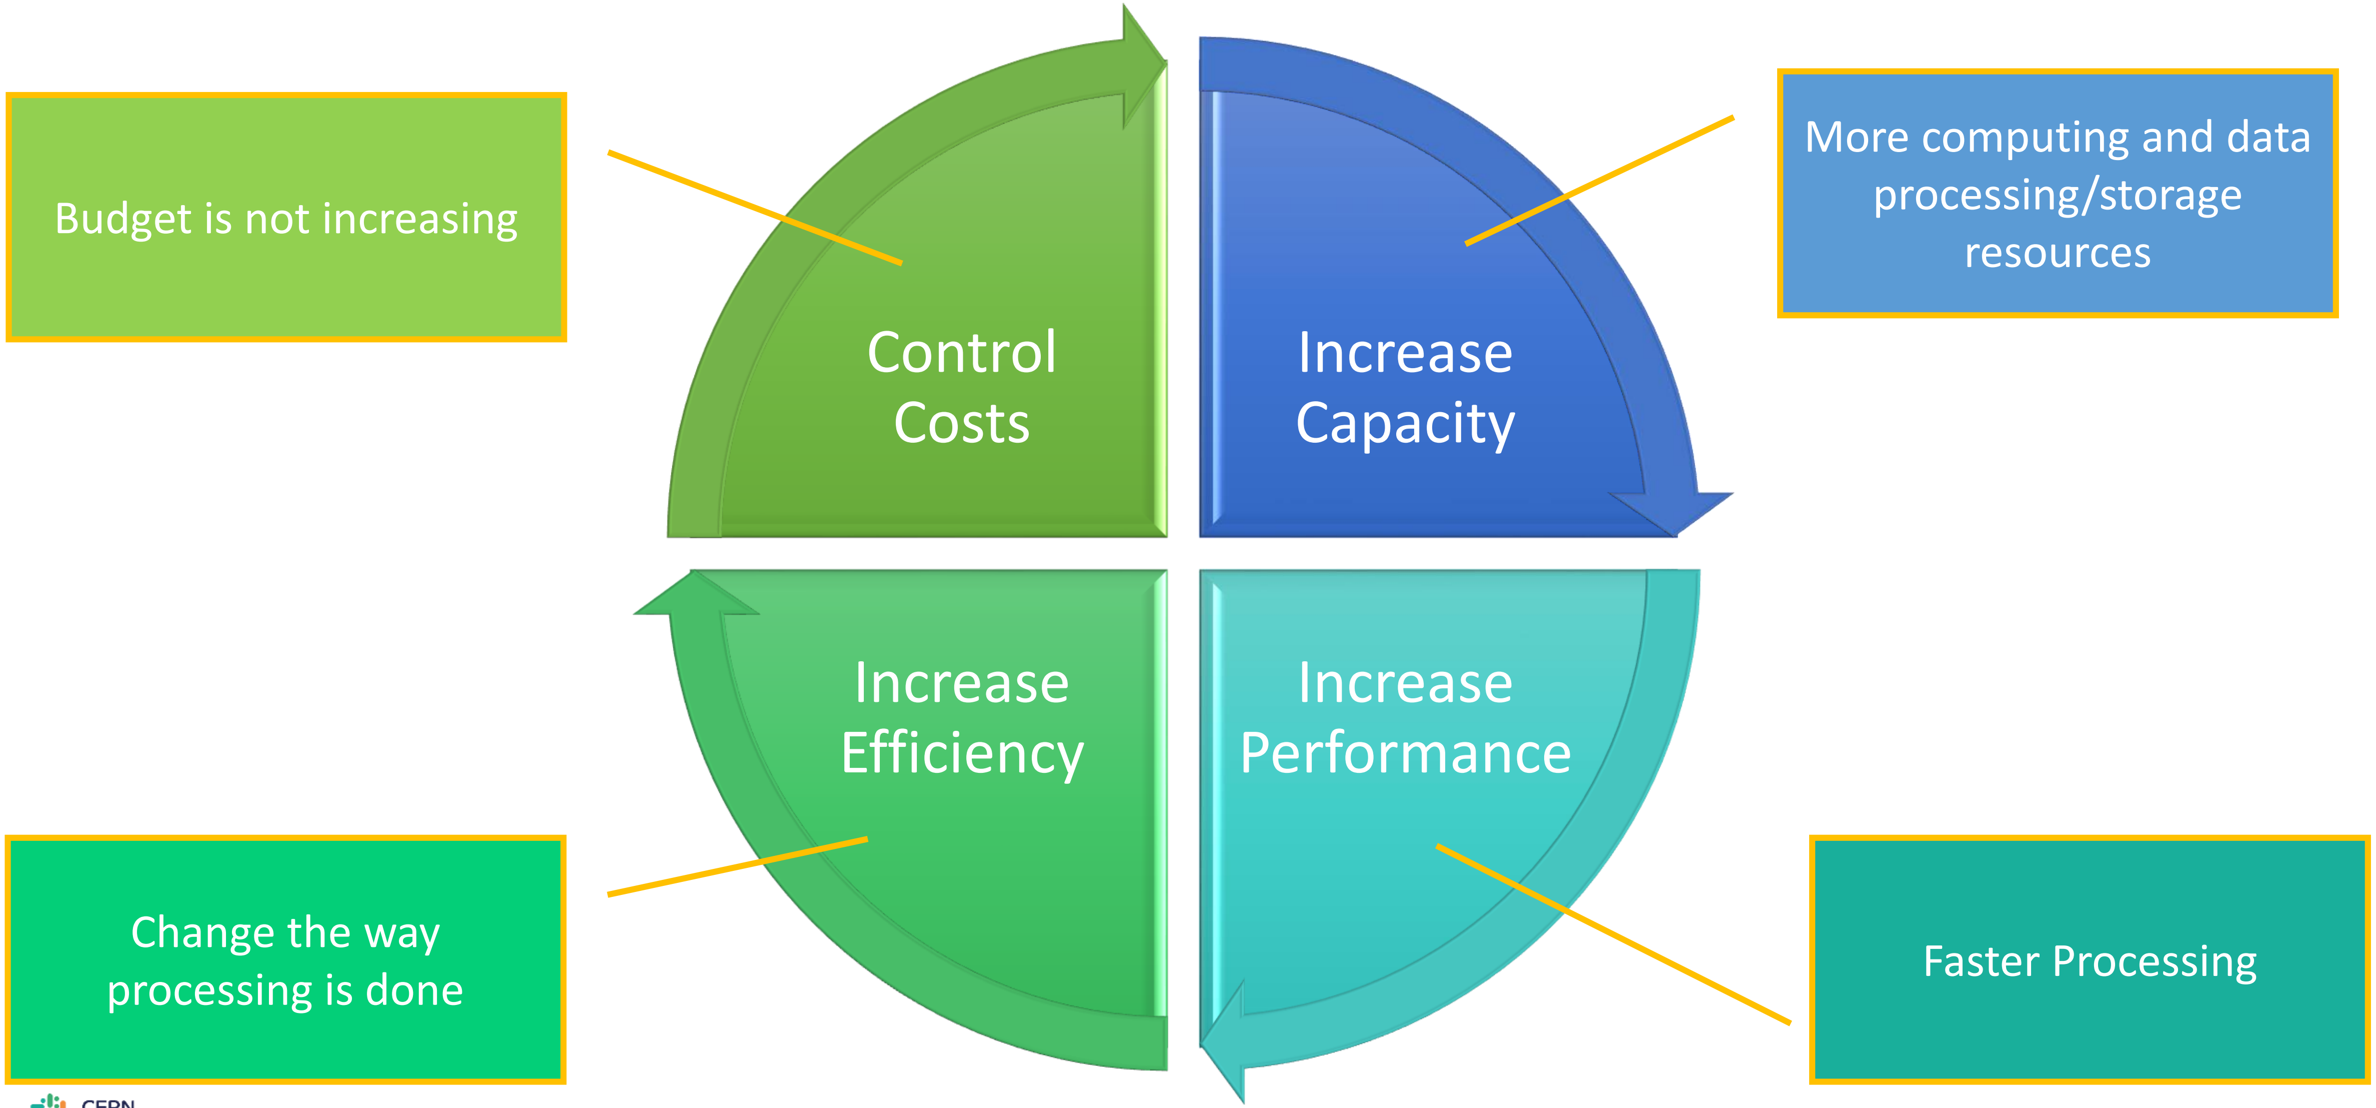

Data Center Architectures

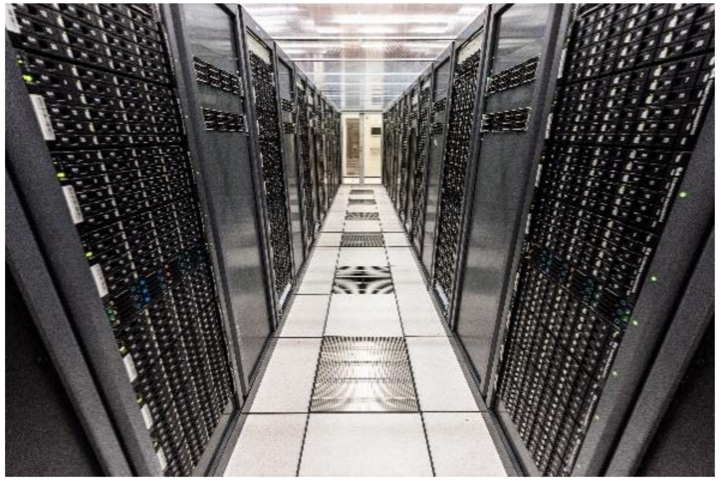

Today traditional data centre architecture Cooled racks, standardized server units, some racks more computing-oriented, some more storage-oriented, 80/20 split between HD/SSD

Investigation of disaggregation (OpenCompute, Intel RSD) as a way of making the use of resources more flexible

Dynamically compose hardware using a software orchestration layer below or integrated with the cloud provisioning layer

Storage Technologies

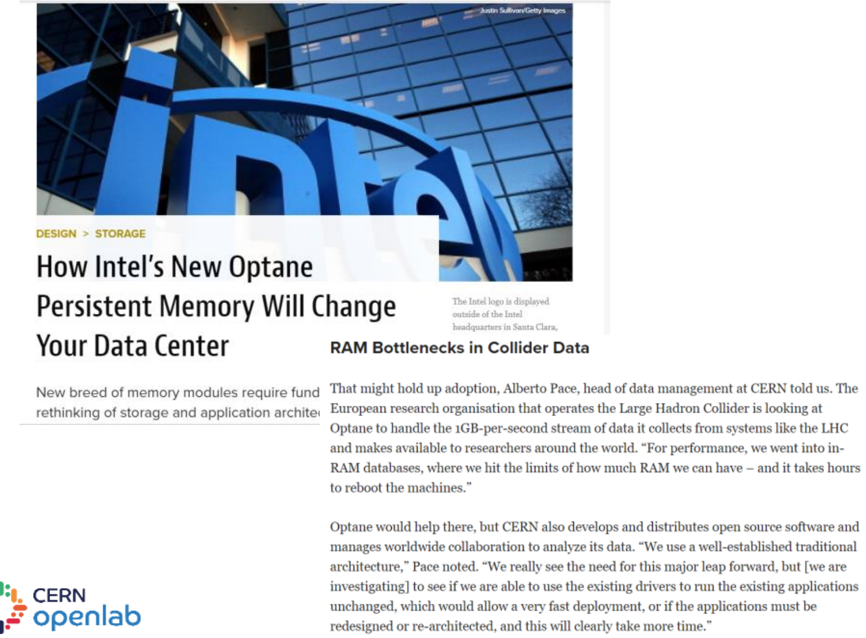

Justin Sullivan/Getty Images

DESIGN > STORAGE

How Intel's New Optane Persistent Memory Will Change Your Data Center

The Intel logo is displayed outside of the Intel headquarters in Santa Clara,

New breed of memory modules require fund rethinking of storage and application archite

RAM Bottlenecks in Collider Data

That might hold up adoption, Alberto Pace, head of data management at CERN told us. The European research organisation that operates the Large Hadron Collider is looking at Optane to handle the 1GB-per-second stream of data it collects from systems like the LHC and makes available to researchers around the world. "For performance, we went into in-RAM databases, where we hit the limits of how much RAM we can have – and it takes hours to reboot the machines."

Optane would help there, but CERN also develops and distributes open source software and manages worldwide collaboration to analyze its data. "We use a well-established traditional architecture," Pace noted. "We really see the need for this major leap forward, but [we are investigating] to see if we are able to use the existing drivers to run the existing applications unchanged, which would allow a very fast deployment, or if the applications must be redesigned or re-architected, and this will clearly take more time."

Storage is a critical component: how do we keep increasing capacity? Can we even do that?

Progressive move towards SSD, but we need very large pools of cheap devices with efficient QoS management for different use cases

Can SSD-based storage replace tapes for long-term storage?

The new 3D-Xpoint devices are very promising for many applications, both as fast disks and for in-memory applications, but cost is a concern

Cloud Computing

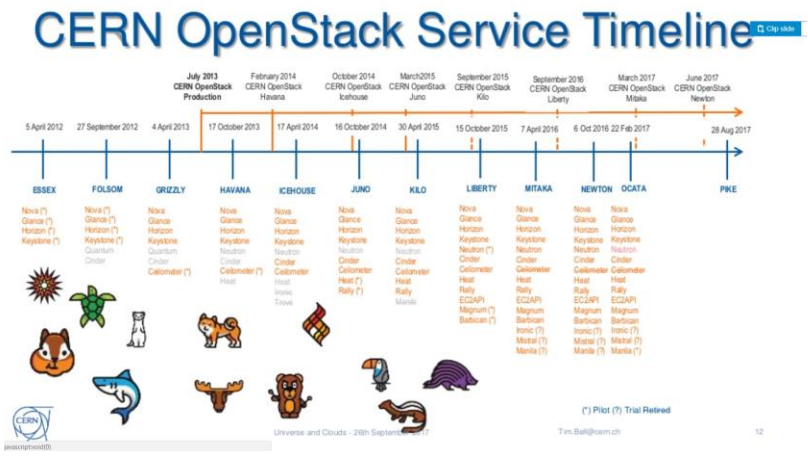

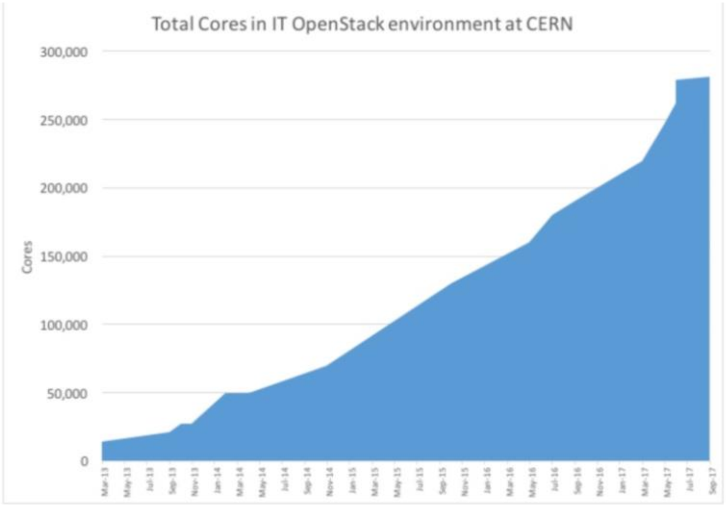

Currently >8000 hypervisors, 281K cores running 33,000 VMs

CERN runs the largest OpenStack-based private cloud in the scientific research world

Use of cloud technology allowed to increase the resource provisioning efficiency by orders of magnitude

However, it is still limited by the maximum physical capacity of the DC

Hybrid public-private models are being considered, also to provision types of resources that we do not have (large GPU farms? Dedicated data training facilities?)

Network Infrastructure

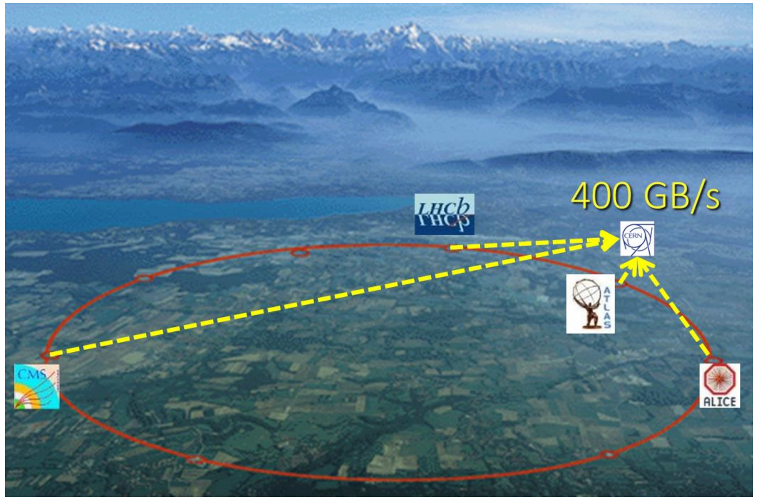

Ethernet, wi-fi and mobile connectivity across two CERN sites for up to 15000 users working at surface and underground installations

Separate public, technical and WLCG operations networks, an IXP node for international telecom providers

Investigating the impact of IoT and the potential applications of 5G networks

To centralize some of the LHC Experiments post-DAQ workloads, bandwidth must increase to multi TB/s over 10s of km (C/DWDM?)

New Computing Platforms

Moore's Law

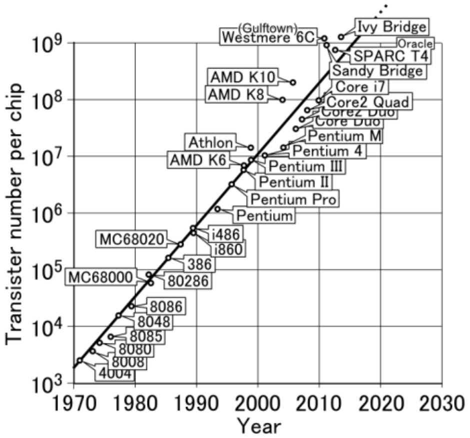

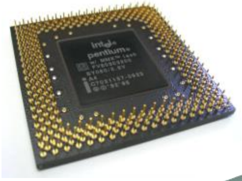

1995

2000

2010

2014

2015

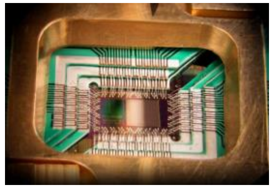

… and beyond

Radically new computing platforms are rapidly moving from pure computer science to realistic devices, e.g. **Neuromorphic Computing** and **Quantum Computing**

A **Quantum Computing Initiative** has been launched in December 2018 as a long-term investigation activity

Code Modernization

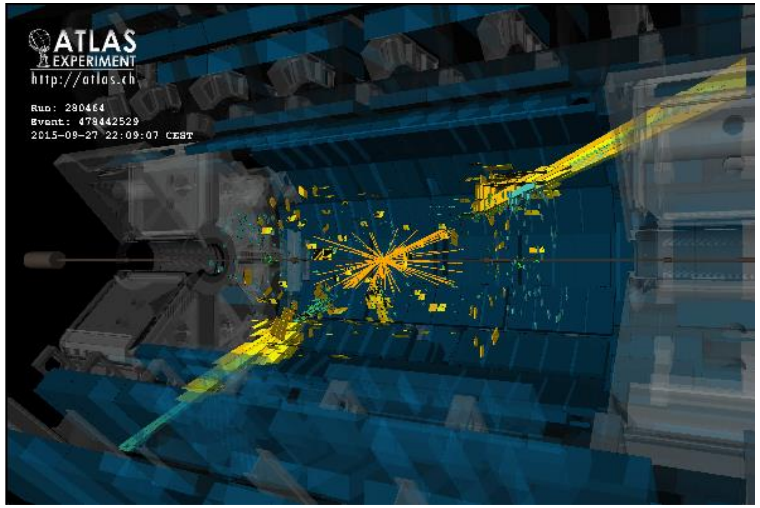

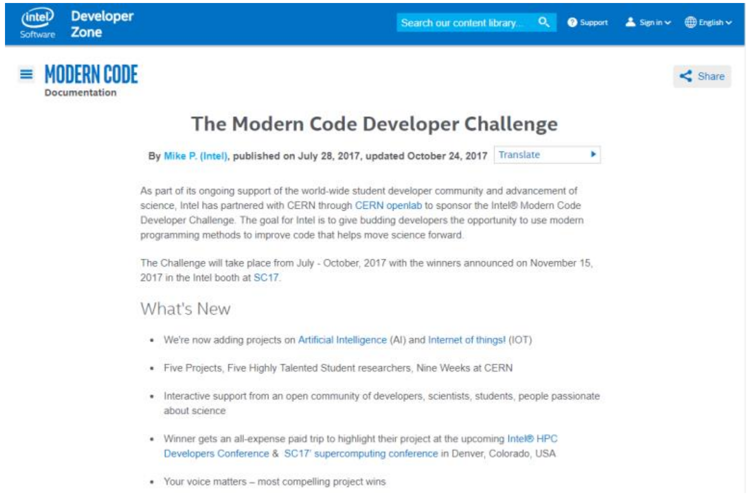

Developer Zone
Search our content library...
Support
Sign in
English

MODERN CODE
Documentation

Share

The Modern Code Developer Challenge

By Mike P. (Intel), published on July 28, 2017, updated October 24, 2017 Translate

As part of its ongoing support of the world-wide student developer community and advancement of science, Intel has partnered with CERN through CERN openlab to sponsor the Intel® Modern Code Developer Challenge. The goal for Intel is to give budding developers the opportunity to use modern programming methods to improve code that helps move science forward.

The Challenge will take place from July - October, 2017 with the winners announced on November 15, 2017 in the Intel booth at SC17.

What's New

- We're now adding projects on Artificial Intelligence (AI) and Internet of things! (IOT)
- Five Projects, Five Highly Talented Student researchers, Nine Weeks at CERN
- Interactive support from an open community of developers, scientists, students, people passionate about science
- Winner gets an all-expense paid trip to highlight their project at the upcoming Intel® HPC Developers Conference & SC17' supercomputing conference in Denver, Colorado, USA
- Your voice matters – most compelling project wins

Most of the software used in HEP was written many years ago, it is mostly single-threaded

A thorough revision of the most important packages has started to introduce vectorization and parallelization as extensively as possible

CERN is an IPCC grant recipient with pioneering work on “modernizing” large packages like Geant, the primary physics simulation tool

We run “code modernization challenges” with Intel where students worldwide are given code and tools and compete for the best optimizations

Data Analysis

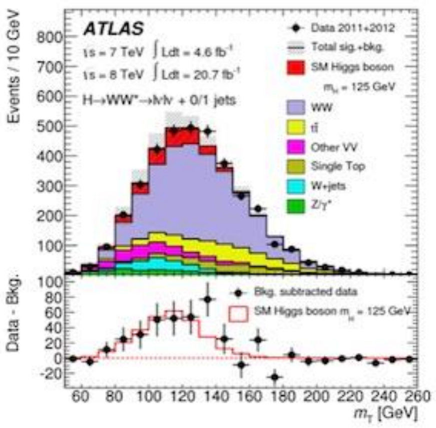

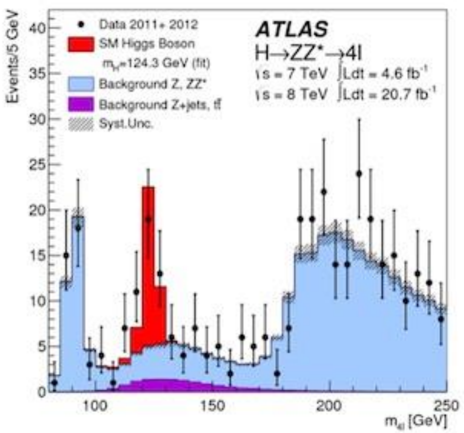

Data analysis is HEP bread and butter, hundreds of PB of data are continuously processed across a worldwide infrastructure

The standard analysis tool today is Root, a data framework designed to handle very large quality of data and provide all necessary management and visualization tools

How do we compare to recent trends everywhere else in the world? Are tools like Hadoop or Spark more or less efficient, more or less cost-effective?

ML/DL for Data Acquisition

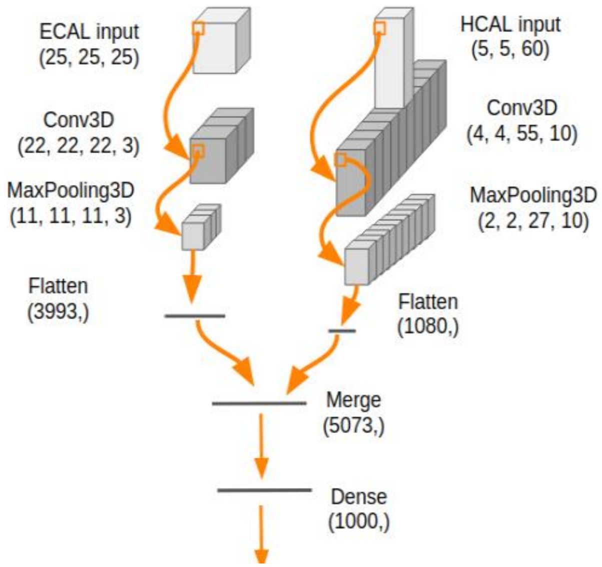

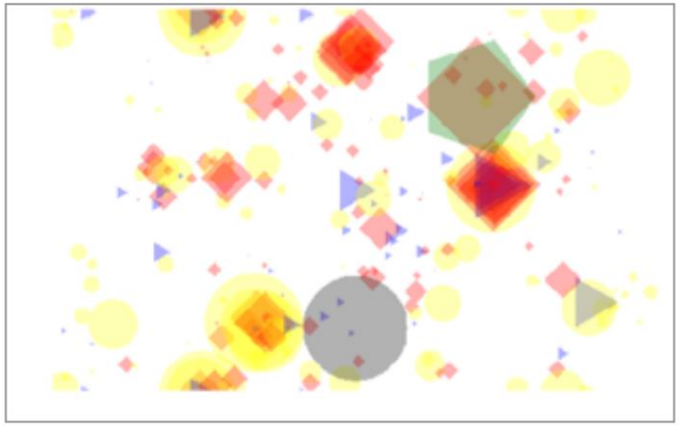

Images courtesy of Maurizio Pierini, CERN CMS

LHC Experiments have used ML for quite some time now (e.g. BDT), now investigating DL in several areas

As collision rates move from 40 MHz to 200 MHz, fast event classification is critical

From convolutional NN for event image (tracks) analysis to NLP-like approach to for “reading” the detector output with the Physics models as the grammar of a language of nature

Initial tests show potential for performance gains of up to 10^3

ML/DL for Simulation

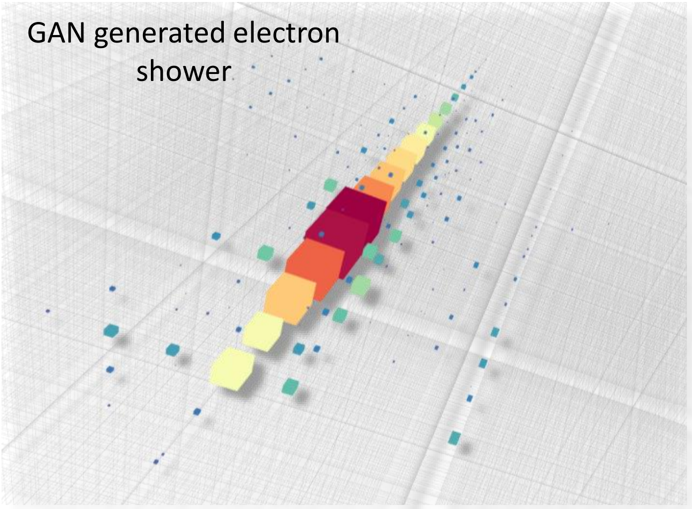

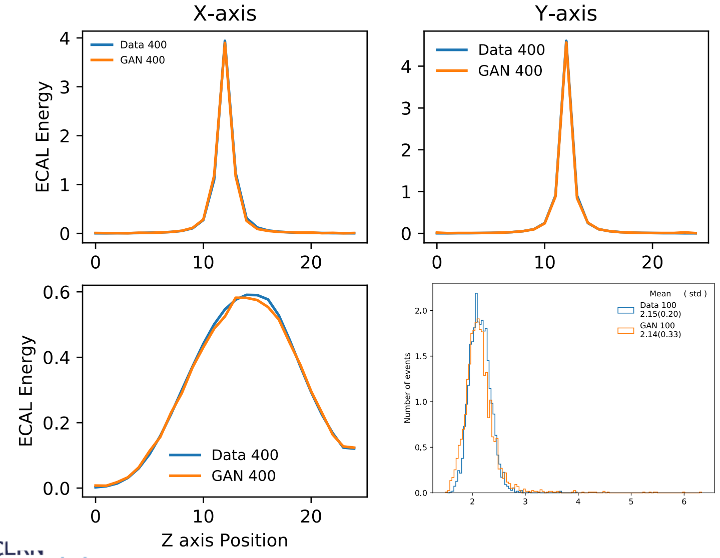

Simulation is a very big part of the HEP data analysis process. Traditionally done with Monte Carlo methods.

Also in this case serious concerns about the scalability of MC in the HL-LHC era

Currently investigating the use of GAN methods to speed up simulation

Inference has already been demonstrated to be at least x10^3 faster, training still an issue

Images courtesy of Sofia Vallecorsa, CERN openlab

ML/DL for Control Systems

CERN manages the Accelerator Complex, a sophisticated machine made of many different accelerators and support subsystems (power, cryogenics, cooling/ventilation, etc.)

The control systems include more than 30 millions sensors and produce continuous stream of TBs of data from edge to central DBs

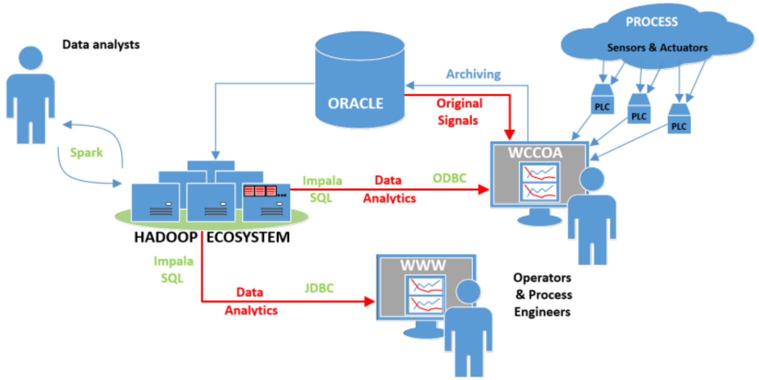

ML is used to make sense of this data, but the problem is to make it distributed through the sensor network

Predictive maintenance and self-healing systems are the objectives

Images courtesy of Filippo Tilaro and CERN, CERN BE

ML/DL for Infrastructure Monitoring

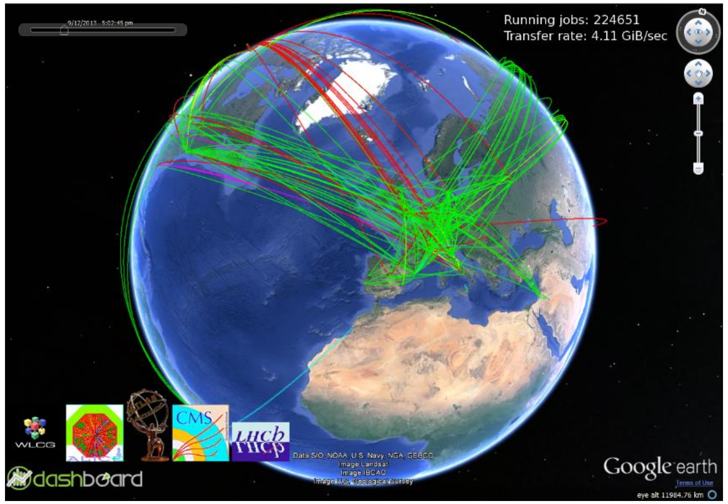

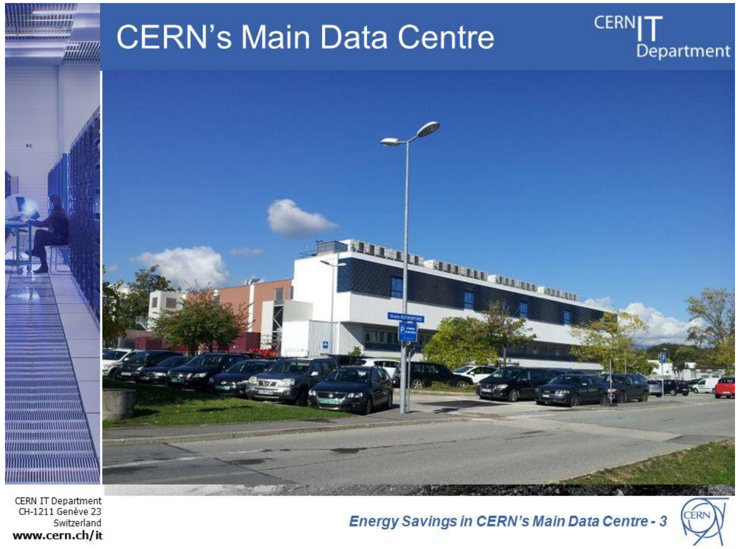

Application of ML techniques are also being evaluated to optimize several infrastructure tasks

Data placement: use smart data analysis to predict where to move data across the WLCG infrastructure

Network security: analyse traffic patterns to detect anomalies and intrusions

Data Centre optimization: analyse FRUs data to optimize job allocation, resource utilization, energy consumption, etc.

IP Management, Open*

- The basic principle of any CERN openlab collaboration is openness
- We assume shared IP of results among project members
  - More specific IP agreements can be discussed with the CERN KT Office
- Within the respect of limited confidentiality agreements and short embargo periods, we expect the results of the projects to be released to the scientific communities following open policies

open source

open hardware

OPEN ACCESS

KNOWLEDGE SHARING

Working with communities beyond high-energy physics

Working closely with CERN KT on initiatives aimed at transferring tools, skills, and knowledge from the high-energy physics community to other research fields.

Aligned with the recently approved knowledge transfer strategy for the benefit of medical applications

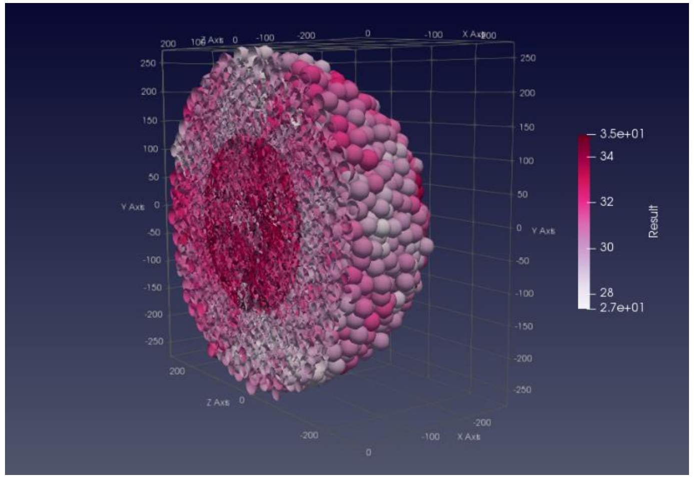

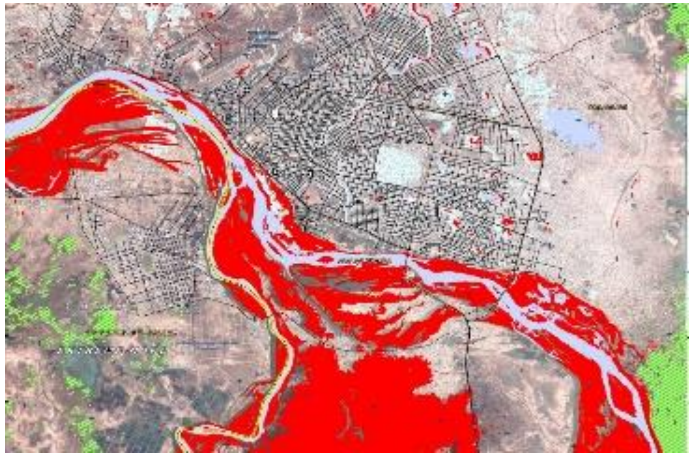

“SMART THINGS” AND IOT

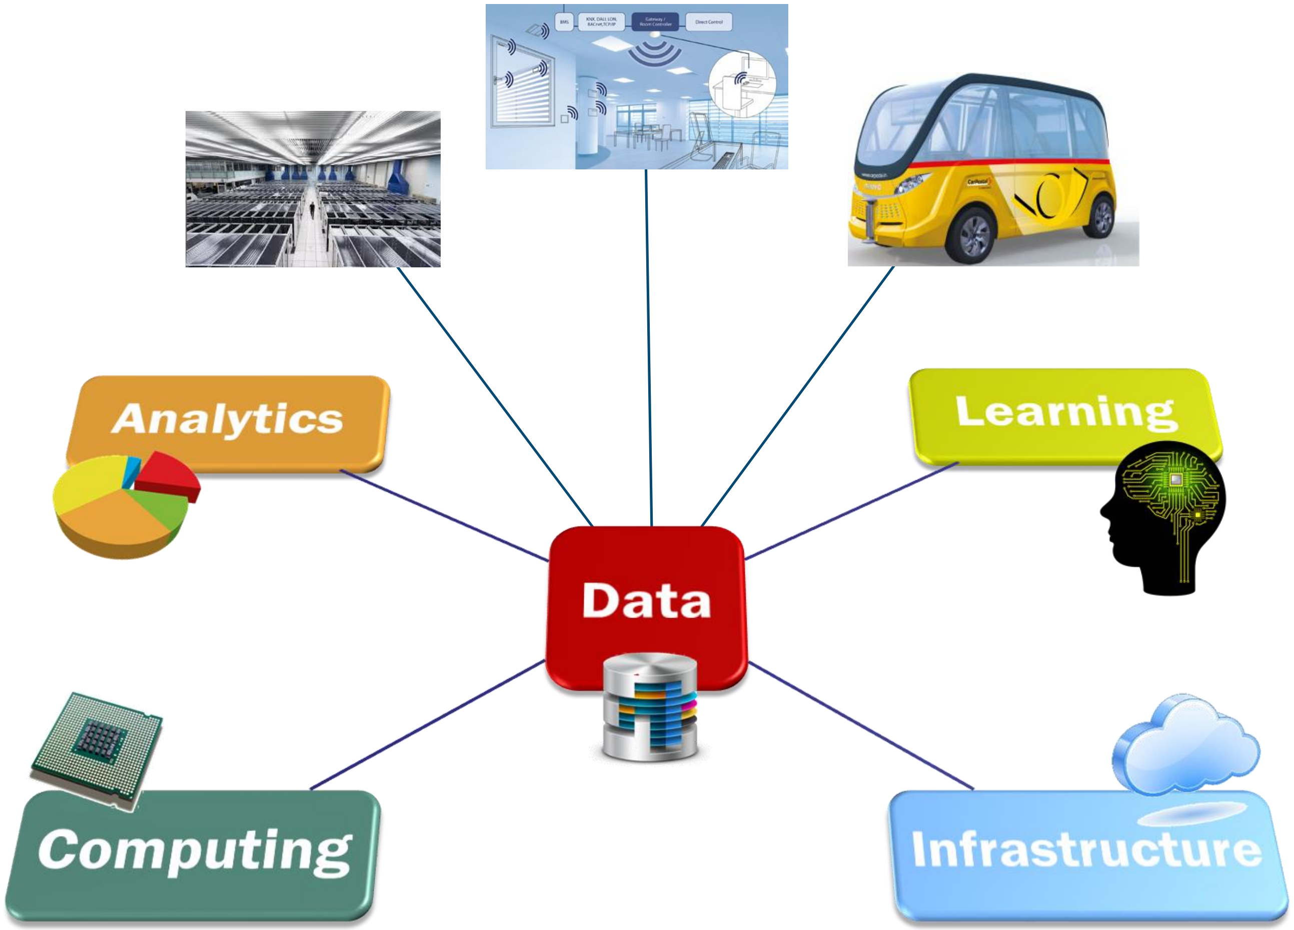

Understand the potential and impact of technologies such as Internet of Things, fast wireless/mobile communication (5G), and large-scale DA/ML

Raise awareness among the community by collecting and aggregating interests and setting up PoC projects

Main interests so far: data center operations optimization, environmental control, mobility

CERN is well-placed geographically and technologically to play a central role

Current status: Requirements analysis

Seeking collaborations: developers, domain experts, testers, use cases

Smart Knowledge Platforms

- **Common challenge** across many activities: harness the growing amounts of information being produced every second of our lives
- **Convergence** of fast hardware, smart algorithms, and increasingly sophisticated models is getting us closer to practical solutions



Education	Research	Industrial/Social
Adaptative, personalized education environments, guiding the students to achieve their learning objectives	Data Analysis, Preservation, Reproducibility, Knowledge Discovery and Sharing platforms, automating complex tasks, suggesting non-obvious links across disciplines and people	Smart personal assistants informing you about your environment, the use of your personal information, and your rights

CERN openlab

MEASURING IMPACT

Communication and Outreach

- **Our website**: http://openlab.cern/
  ~150k unique visitors/year (+100%)
- Articles in other **CERN and external channels**
  ~900 press cuts (+300%)
- **Facebook**: http://cern.ch/go/p7pF
  7.3k people reached / 100k single posts
- **Facebook group**: http://cern.ch/go/n6xD
- **Twitter**: http://cern.ch/go/7vvk
- **Workplace**: http://cern.ch/go/6MTQ
- **LinkedIn**: http://cern.ch/go/NK9k
- Join the **alumni**: https://alumni.cern/

EDUCATION AND TRAINING

Training tomorrow's leaders in ICT

CERN openlab runs a ***highly competitive*** summer-student programme

~ 40 students selected from around 1700 applicants

Most dedicated CERN openlab personnel are ***young, talented 'fellows'.***

Receive hands-on experience with cutting-edge technologies

Workshops, training courses, and other initiatives run with our collaborators.

Competitions, challenges, online training, hackathons.

Experts from industry and research give lectures.

Participate in events both inside and outside CERN.

SUMMER STUDENT PROGRAMME

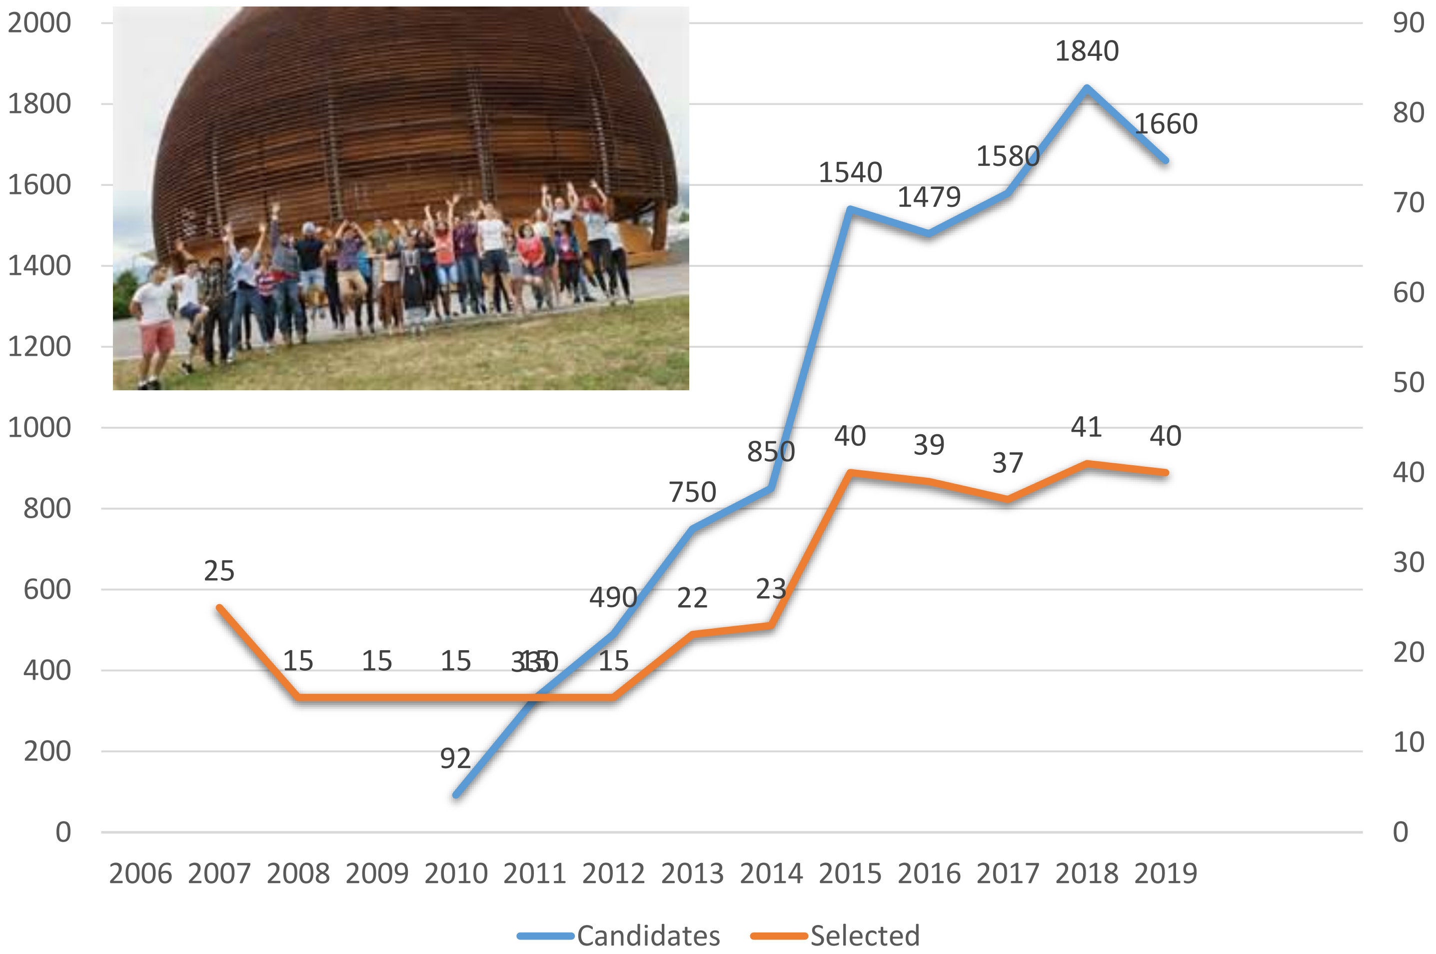

In 2019

- 1660 applicants
- 40 selected students
- 14 lectures
- Visits to external labs and companies
- Lightning talks session
- Technical reports

CONTACTS

ALBERTO DI MEGLIO
CERN openlab Head
alberto.di.meglio@cern.ch

MARIA GIRONE
CERN openlab CTO
maria.girone@cern.ch

FONS RADEMAKERS
CERN openlab CRO
fons.rademakers@cern.ch

FEDERICO CARMINATI
CERN openlab CIO
federico.carminati@cern.ch

ANDREW PURCELL
CERN openlab Communications Officer
andrew.purcell@cern.ch

KRISTINA GUNNE
CERN openlab Administration/Finance Officer
kristina.gunne@cern.ch

www.cern.ch/openlab

Thanks!

alberto.di.meglio@cern.ch
@AlbertoDiMeglio