



Phase VII and Beyond

Alberto Di Meglio – CERN openlab Head

11/03/2021

A Time of Transition and Planning



CERN openlab
Phase VI

CERN openlab
Phase VII

20th Anniversary
and beyond

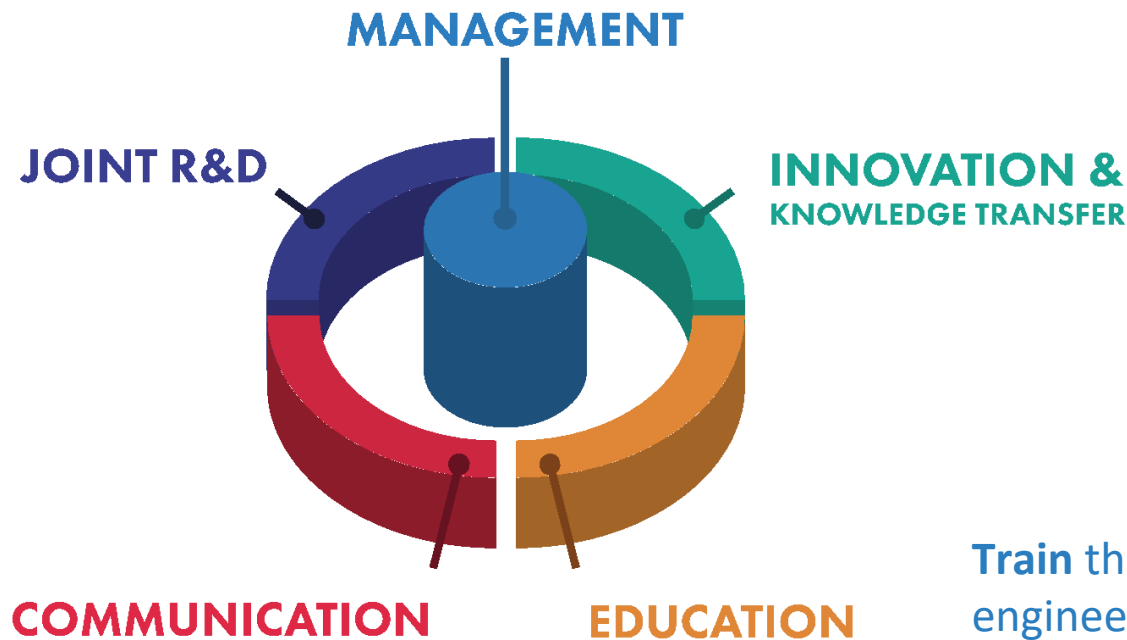
CERN openlab Phase VII and Beyond

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.

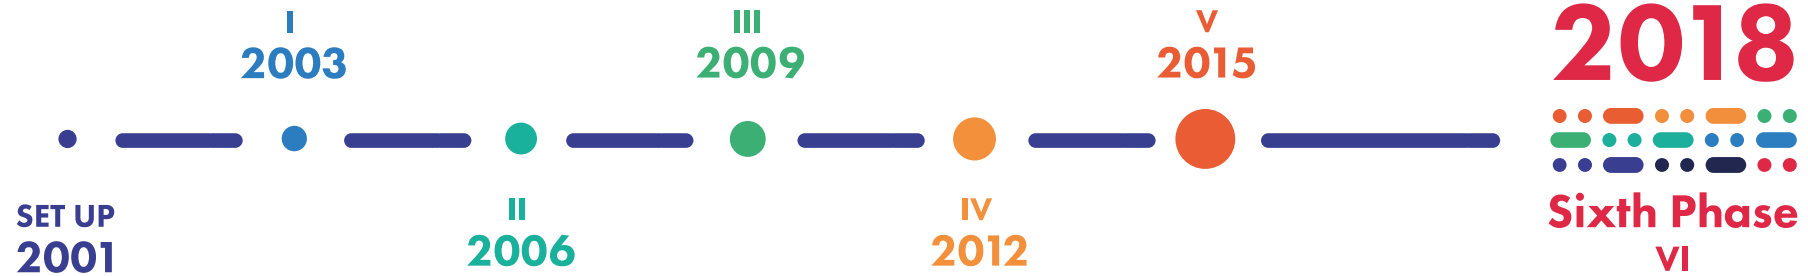
Communicate results, demonstrate impact, and reach new audiences.



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

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

DRIVING INNOVATION SINCE 2001



JOINT R&D PROJECTS (PHASE VI)



High-bandwidth fabrics,
accelerated platforms for
data acquisition



HPC, Cloud,
Quantum



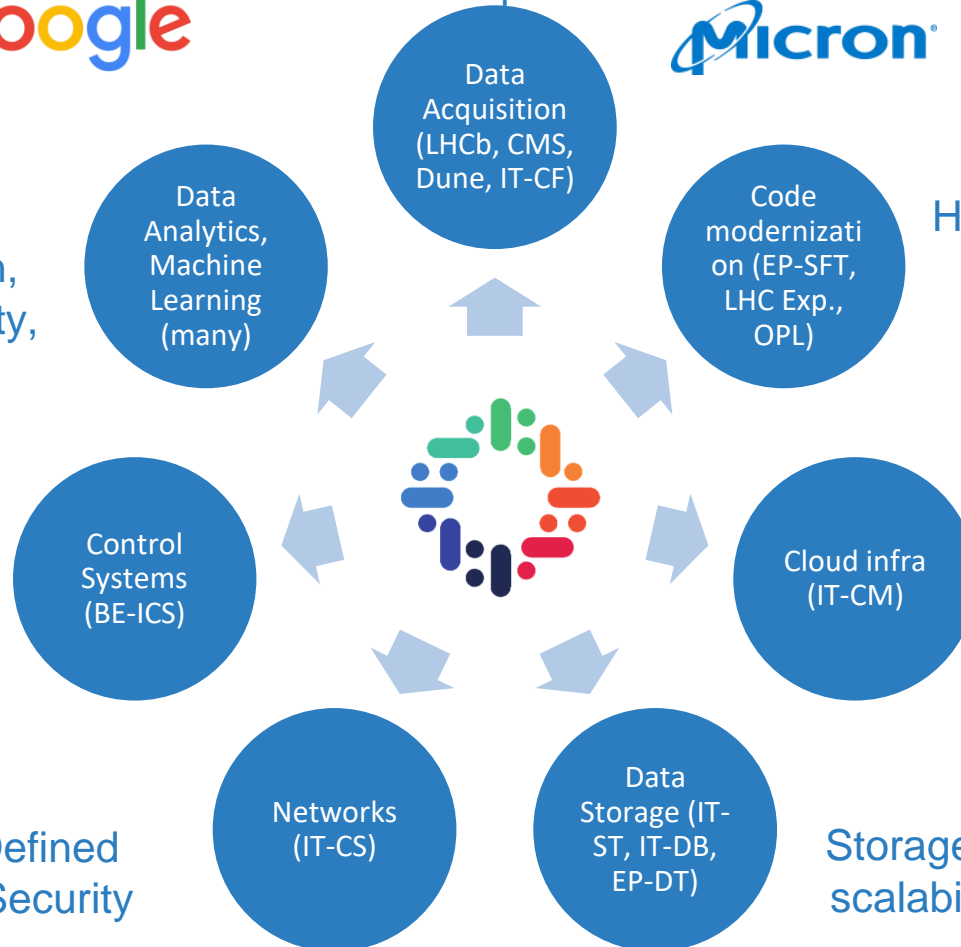
САМАРСКИЙ УНИВЕРСИТЕТ
SAMARA UNIVERSITY

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

Predictive/proactive
maintenance and
operations



Software Defined
Networks, Security



CERN openlab Phase VII and Beyond



Cloud technology,
containers, scalability



Storage architectures,
scalability, monitoring



CERN openlab VI in Numbers

Phase VI was a very successful phase!

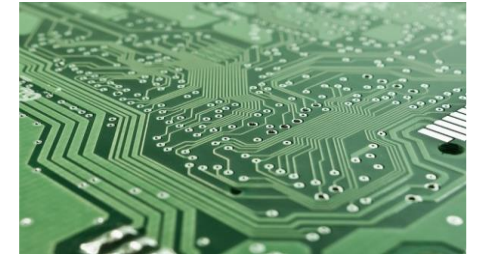
- 16 industry members
- 10 academic/research/no-profit members
- ~35 projects over 3 years
- Grants for 10 TECH/PJAS, 4 DOCT, 18 Fellows, 5 Staff and ~100 Summer Students/Interns!
- > 100k visitors to website
- > 5k applications to summer-student programme
- ~ 300 press articles
- > 100 main communications (by CERN and members)
- Over 17k followers on our main social media channels.
- CERN openlab channel in the CERN's alumni platform
 - Please join us on there: simply go to **alumni.cern**, sign up, and then request to join the CERN openlab group. Would like to start growing this channel with upcoming 20th anniversary in mind.

Phase VI: Three Main Areas of R&D

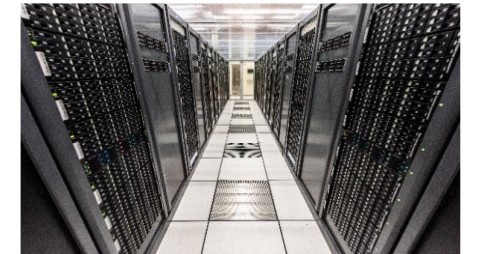


COMPUTING
CHALLENGES

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**



Main Achievements

Data Centre Technologies and Performance

- Evaluation of heterogeneous architectures (Intel, IBM, E4 on NVIDIA architectures, Micron, EU funded project DEEP-EST)
 - GPUs adoption for HLT reconstruction (Allen, Patatrack, DEEP-EST)
 - GPUs for training and inference in ML applications
 - FPGA-based DL with Micron DLA in CMS and protoDUNE
- Unified programming models: investigations with Intel OneAPI
 - Heterogeneous hardware support
- Evaluation of performance
 - Benchmarking suite on HPC
 - Profiling code on multiple architectures
 - Testbed access on multiple architectures
- Cloud solutions for data analytics, database solutions, machine learning, disaster recovery and compute batch, investigation on new challenges brought by Kubernetes WebLogic deployment with Oracle
- Storage solutions
 - Distributed Asynchronous Object Storage (DAOS) open source investigated with Intel
 - EOS productisation with COMTRADE

Main Achievements

Scale out capacity

Most of the activities in this area have focused on investigating how to scale out of premise for ML/DL workloads in terms of both increasing the amount of resources and accessing new types of architectures. For example:

- ML/DL training and data analysis at scale with Kubernetes on Google Clouds
- Large-scale 3D-GAN simulation with Intel architectures at SURFsara
- Satellite image processing, segmentation and generation with UNOSAT on T-Systems OTC
- Initial strategic assessments of HPCaaS, MLaaS and QCaaS on multiple clouds in collaboration with the CloudBank EU project using AWS and GCP resources

Main Achievements

New Computing Paradigms

The goal of this area was the exploration of ML/DL algorithms and their optimisation on new specialised architectures. For example:

- New algorithms for neutrino experiments data processing with Dune and IBM
- Efficient use of new hardware architectures:
 - DL inference acceleration of 2D-GANs through low precision data representation with Intel
 - Optimisation of distributed training on TPUs with Google
- Advanced algorithms:
 - 3D-GAN and Boosted 3D-GAN for calorimeters simulation
 - Progressive GANs for synthetic images generation
 - Key-Point Detection algorithm for Noisy Data
- Establishment of first Quantum Machine Learning pilots for HEP applications
 - qGAN, qGNN, qSVM, hybrid models
- Spin-off of Quantum Computing into the new Quantum Technology Initiative

Phase VII: Recommendations

2020 Update of the European Strategy for Particle Physics

*“the software and computing models used in particle physics research must evolve to meet the future needs of the field” and “the community must **vigorously pursue common, coordinated R&D efforts in collaboration with other fields of science and industry, to develop software and computing infrastructures that exploit recent advances in information technology and data science**”.*

HL-LHC Software and Computing Review Panel Report

Highlights aspects such as improvement of code performance on hardware accelerator architectures or even the need to converge infrastructure projects to integrate in High Performance Computing (HPC) resources. It also highlights **that the LHC computing model must also consider the evolution of the international computing landscape**, such as the European Open Science Cloud (EOSC)

International HEP Strategy Roadmaps

The ongoing Snowmass process has already massively highlighted the need to focus on more integrated use of **HPC, Clouds, ML/DL tools and frameworks, mainstream data analysis tools, quantum technologies** and more

Phase VII: New Challenges

Phase VII coincides with the LHC Run 3, it's an opportunity to consolidate the investigations started in Phase VI and look forward to HL-LHC

1. **Exascale Technologies** are receiving great attention (and funding). HPC, Clouds and interoperability/portability tools for large-scale heterogeneous architectures will keep playing an increasing role in scientific infrastructures.
How can the HEP community use them effectively?
2. **Artificial Intelligence** (or some of its flavours) is rapidly taking first stage in all data processing applications, but it presents challenges in many directions from large-scale training, to interpretability, up to the ethics of science and technology.
What is its role and impact in HEP research?
3. **Quantum Computing** has recently emerged as a potential future game changer.
Is there a role for it in HEP? How do we build knowledge and expertise and prepare?
4. **Scientific Collaborations** are key enabling elements for innovation and economies of scale and information technologies, computer and data science are common tools of the trade and present common challenges across most scientific research disciplines.
How do we work together, develop scalable common approaches and tools?

Four Pillars of Activity

XT eXascale Technologies

A comprehensive investigation of HPC and Cloud infrastructures, frameworks, tools to support key scientific workloads and applications

AI-S Artificial Intelligence for Science

Analysis and development of algorithms, optimisation for new architectures, interpretability, synergies between Physics and other sciences

QTI-C Quantum Technology Initiative - Computing

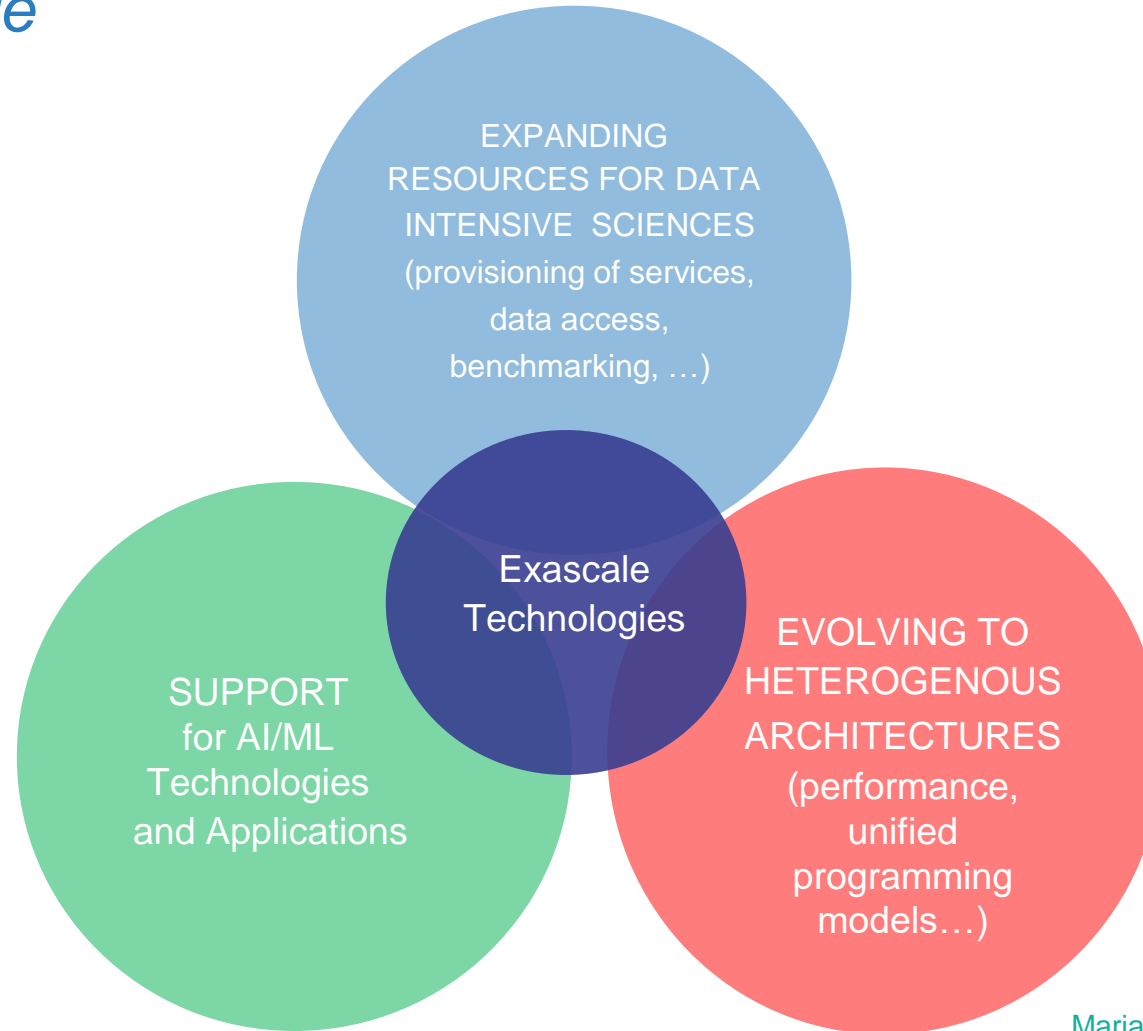
Assess the potential impact of quantum computing in HEP and other sciences, investigate quantum machine learning algorithms and areas of potential quantum advantage, set up a collaborative quantum computing (simulation) platform

MSC Multi-Science Collaborations

Share the expertise and knowledge generated across all activities with other sciences, work with CERN KT to explore novel applications of CERN computing systems and ideas, create collaborations and contribute to common solutions

eXascale Technologies (XT)

HPC, AI and Storage



Maria Girone, CERN openlab Technical Workshop 2021

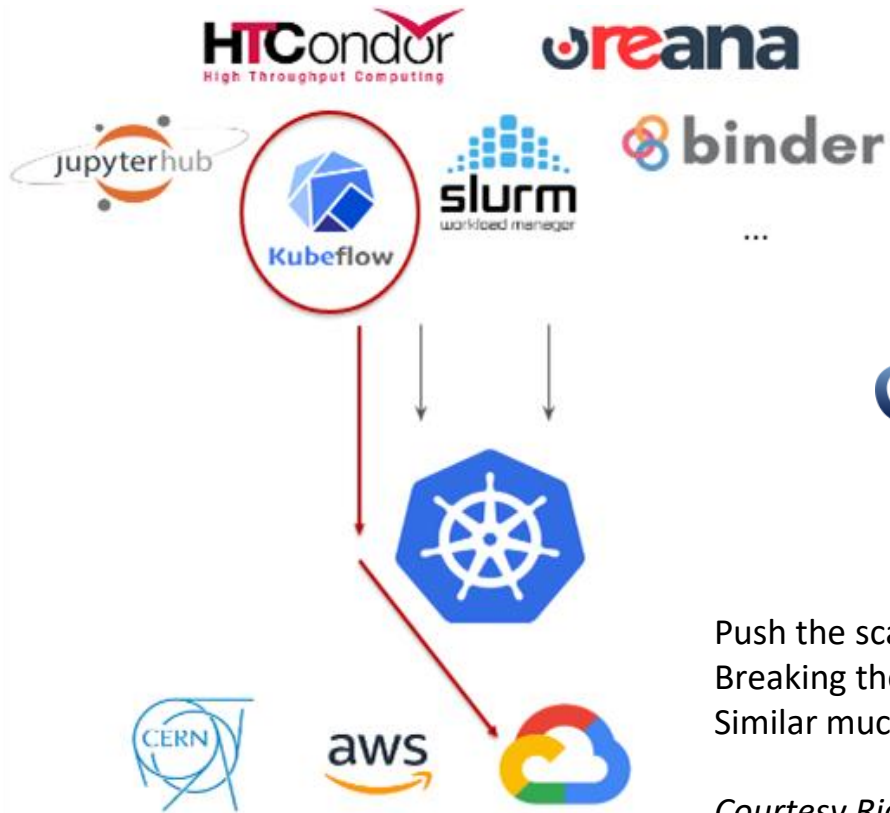
On the path to Exascale

Collaborations



Maria Girone, CERN openlab Technical Workshop 2021

As-a-Service Scale-Out Validation



Push the scale of workloads:
Breaking the 1024 GPU for single workload
Similar much larger TPUs validation tests v3-512 cores

*Courtesy Ricardo Rocha
(research lead IT/ATLAS CloudBank EU project)*



<https://ngiatlantic.eu/news/large-hadron-collider-lhc-farmers-how-society-will-reap-benefits-second-ngiatlanticeu-open>

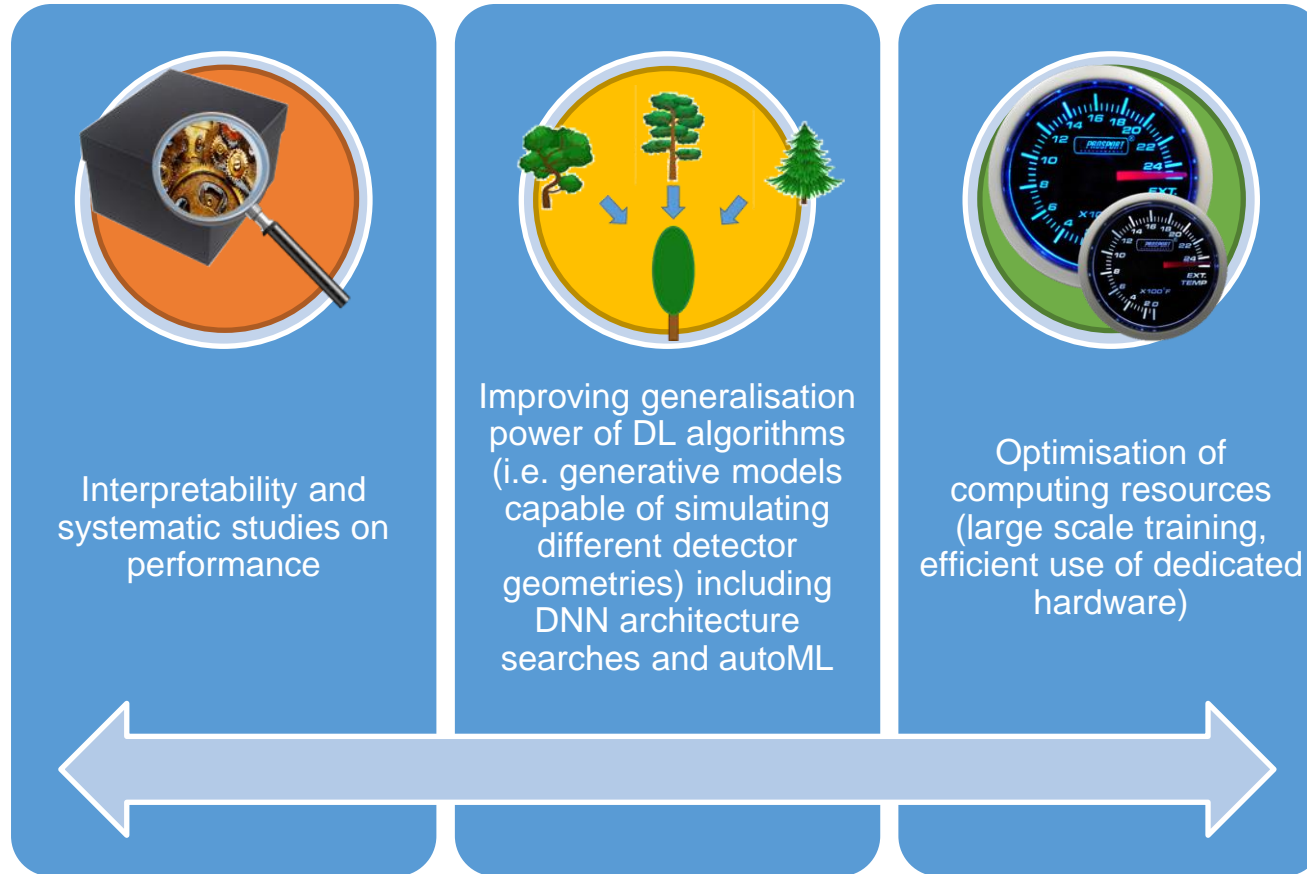
ORACLE®

T-Systems

Microsoft Azure

Artificial Intelligence for Science (AI-S)

Improving usability and trust of ML/DL models



CERN Quantum Technology Initiative (QTI)

Strategy



Joint HEP R&D
Programme



CERN
Management



Advisory Board (**with large
representation of the
Member States**)

Coordination



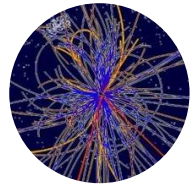
QT Initiative
Management



R&D



Sensing,
Detectors R&D



Computing &
Engineering



Communication



Simulation,
Information
Processing

Capacity building

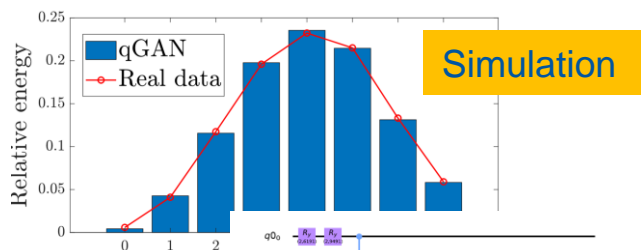
Academic Programmes / Industrial Collaborations



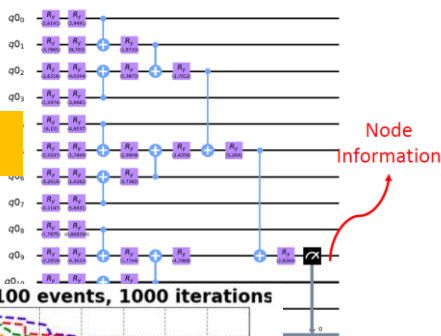
QUANTUM
TECHNOLOGY
INITIATIVE

CERN Quantum Activities

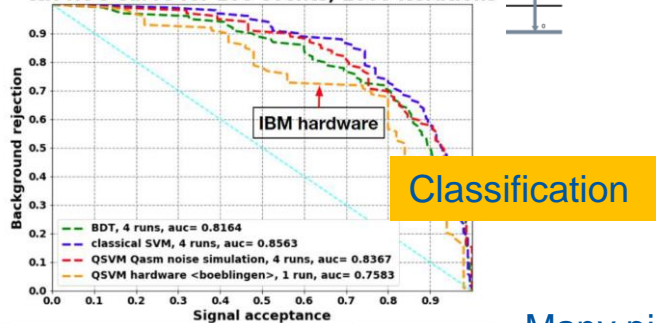
Computing



Reconstruction

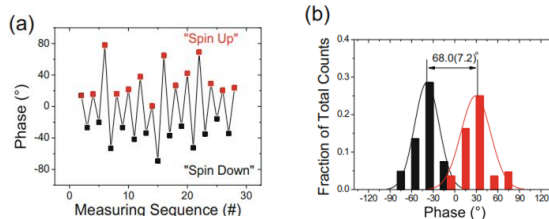


ttH ROC Curve for 100 events, 1000 iterations



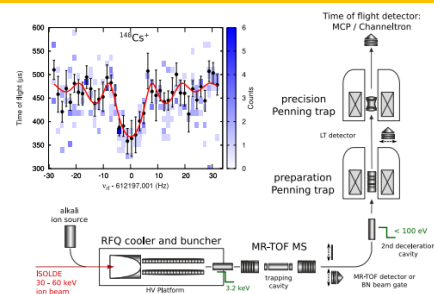
Sensing

BASE - The Baryon Antibaryon Symmetry Experiment



<https://doi.org/10.1140/epjst/e2015-02607-4>

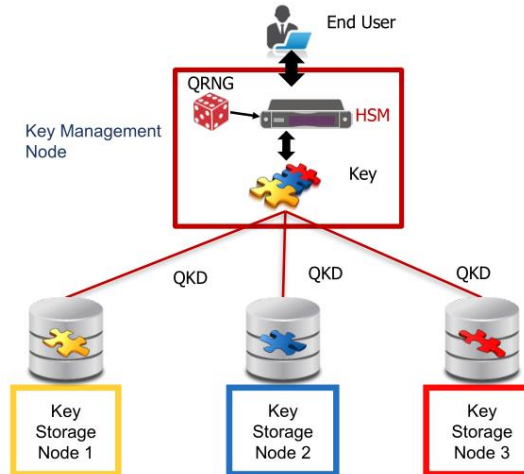
BASE phase-sensitive measurement of spin allowing very precise magnetic field drift measurements



<https://doi.org/10.1088/1361-6471/aa5a20>

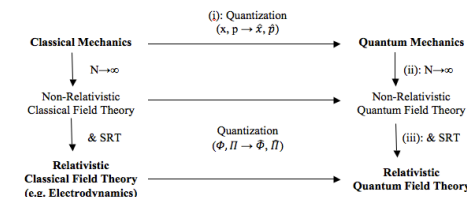
ISOLTRAP Mass-Spec

Communications

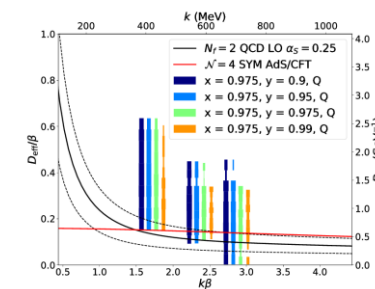


openQKD Repeater node in the CERN Data Centre

Theory



Quantum Field Theory

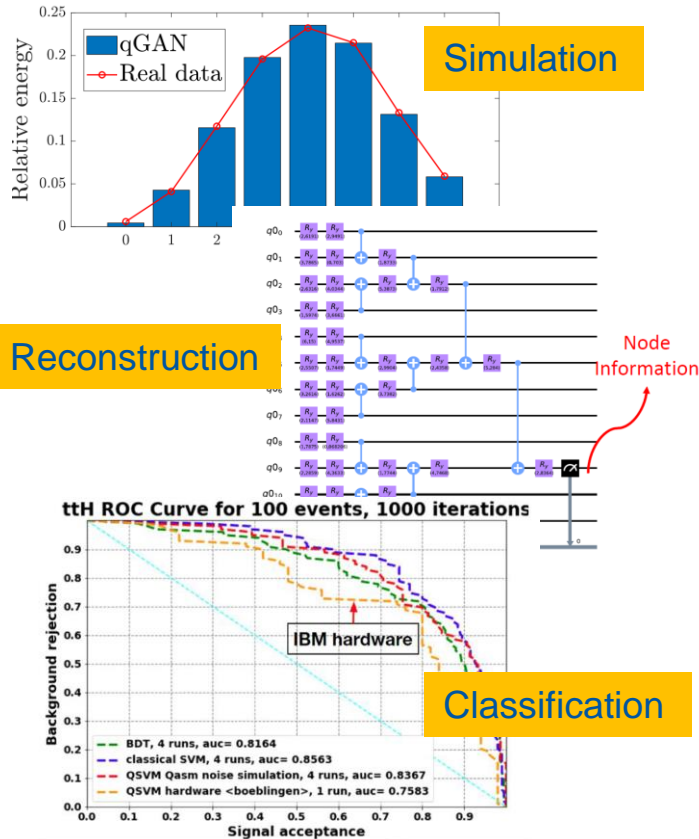


<https://cds.cern.ch/record/2703396>

Lattice QCD

Many pilot projects already started as part of the **CERN openlab quantum** programme (<https://openlab.cern/quantum>)

Quantum Computing

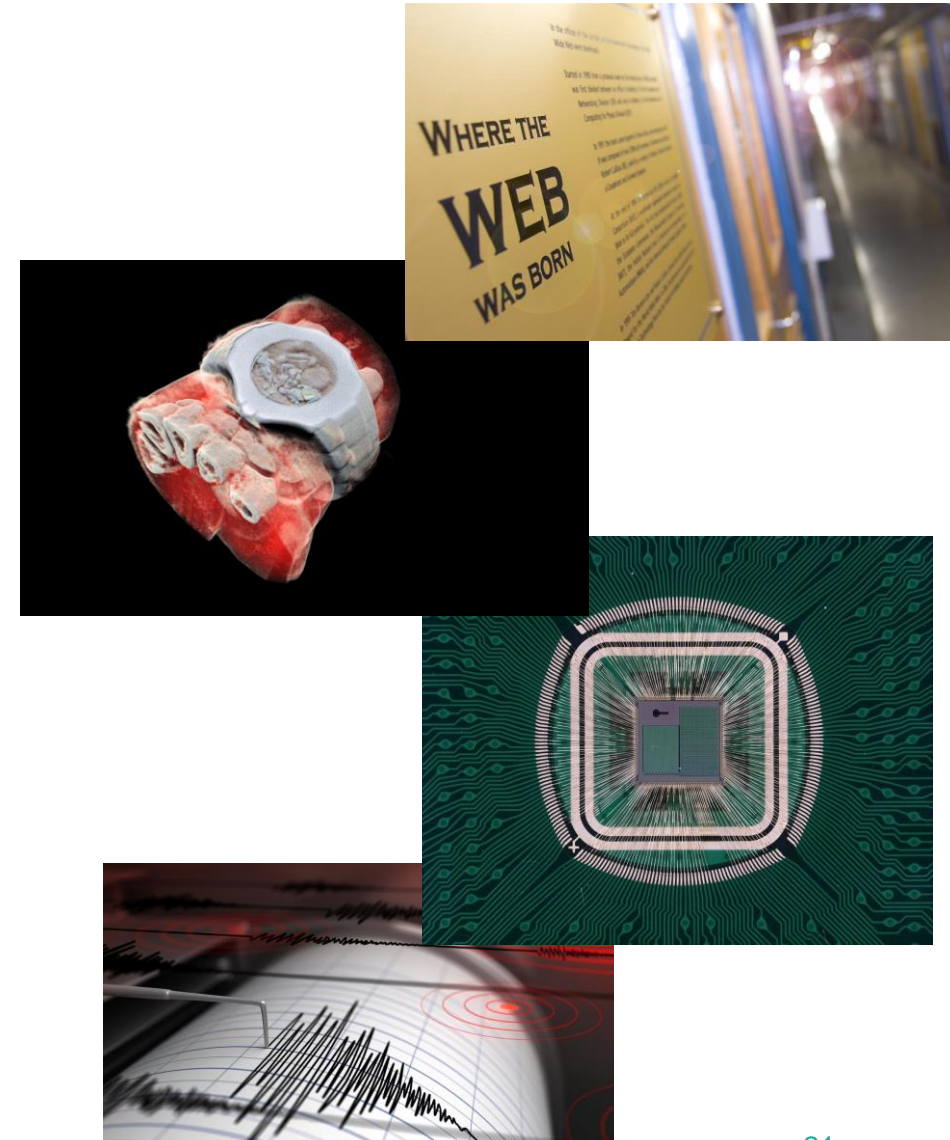
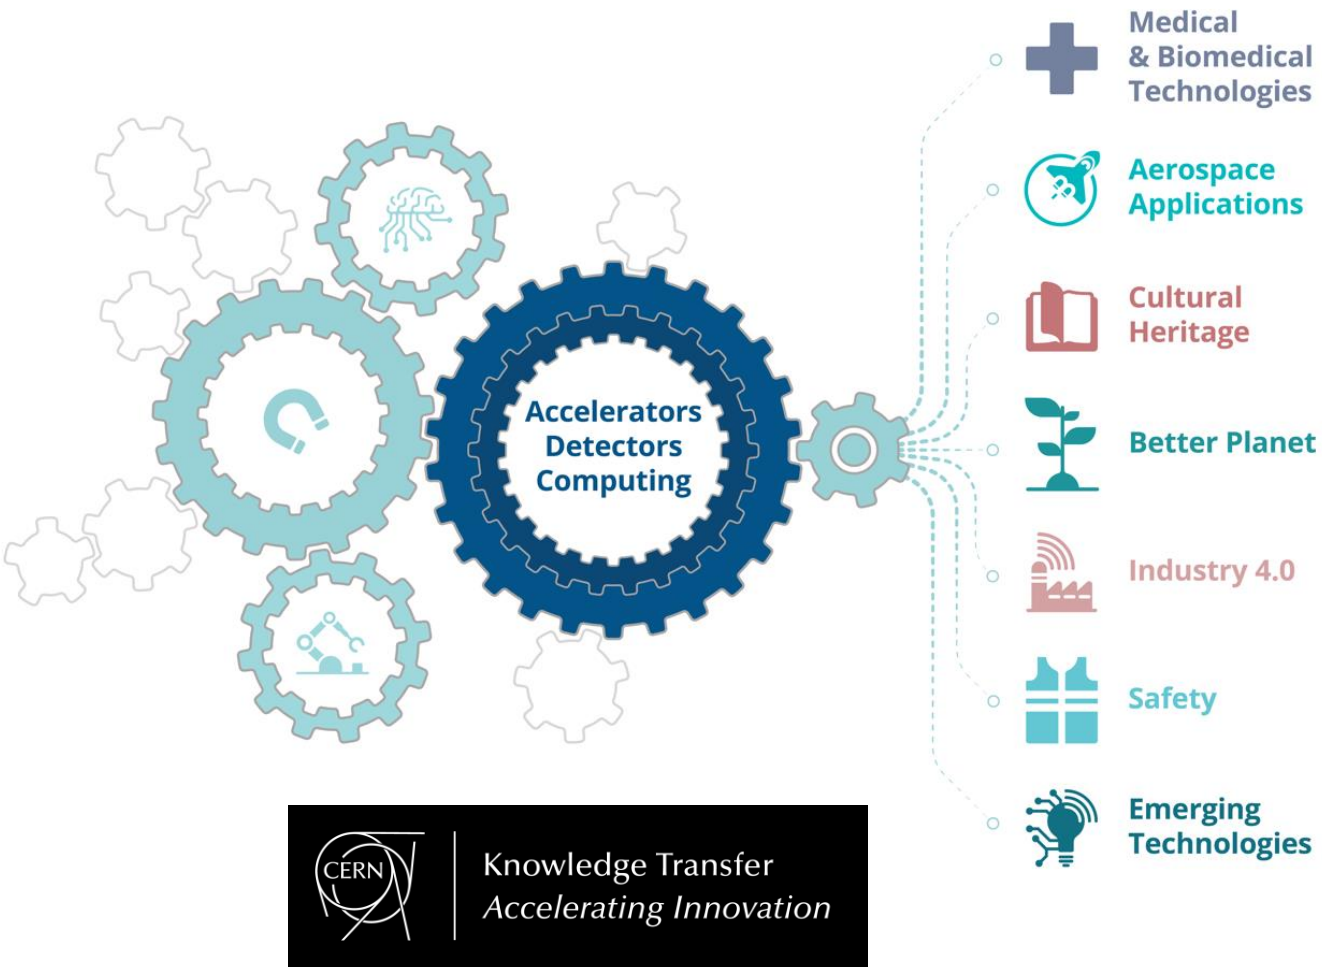


Today: set a baseline for prioritisation and systematisation

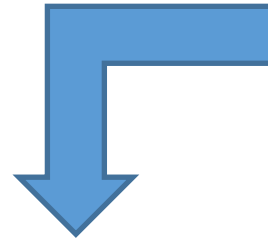
- Quantum **Generative Adversarial Networks** for detector simulation
- Quantum **Graph Neural Networks** for particle trajectory reconstruction
- Quantum **Support Vector Machines** for signal/background classification (Higgs, SUSY,..)
- Workload optimization via quantum **Reinforcement Learning**
- Quantum **Random Number Generators** tests and integration
- Quantum **Homomorphic Encryption**

Later: focus on a more formal approach to algorithms, methods, error characterisation and correction

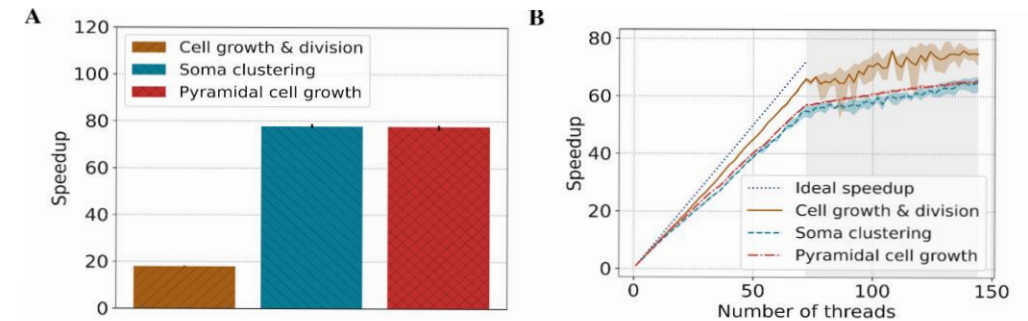
Knowledge Sharing and Transfer



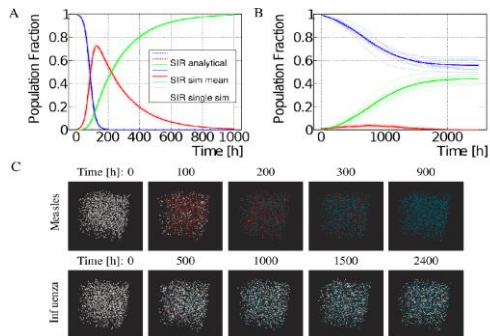
BioDynaMo



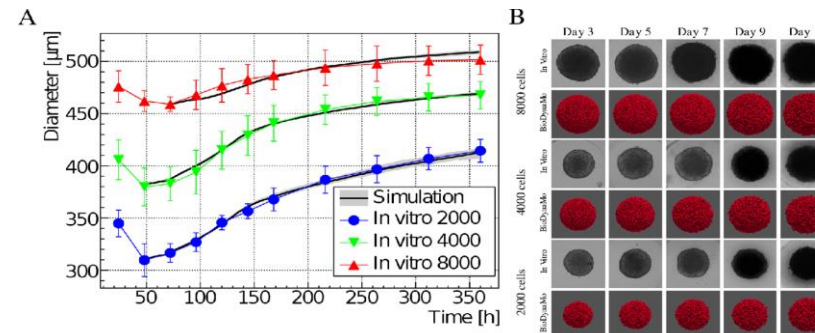
High-performance simulation platform



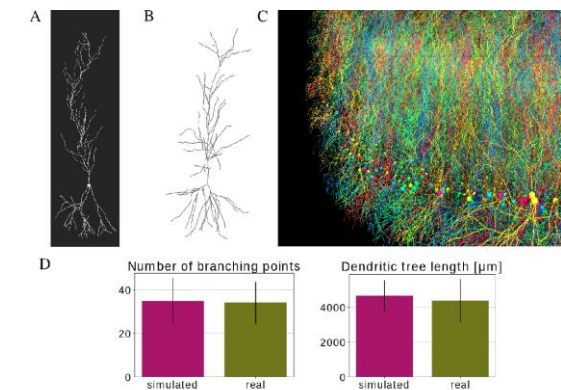
Epidemiology use case



Oncology use case



Neuroscience use case



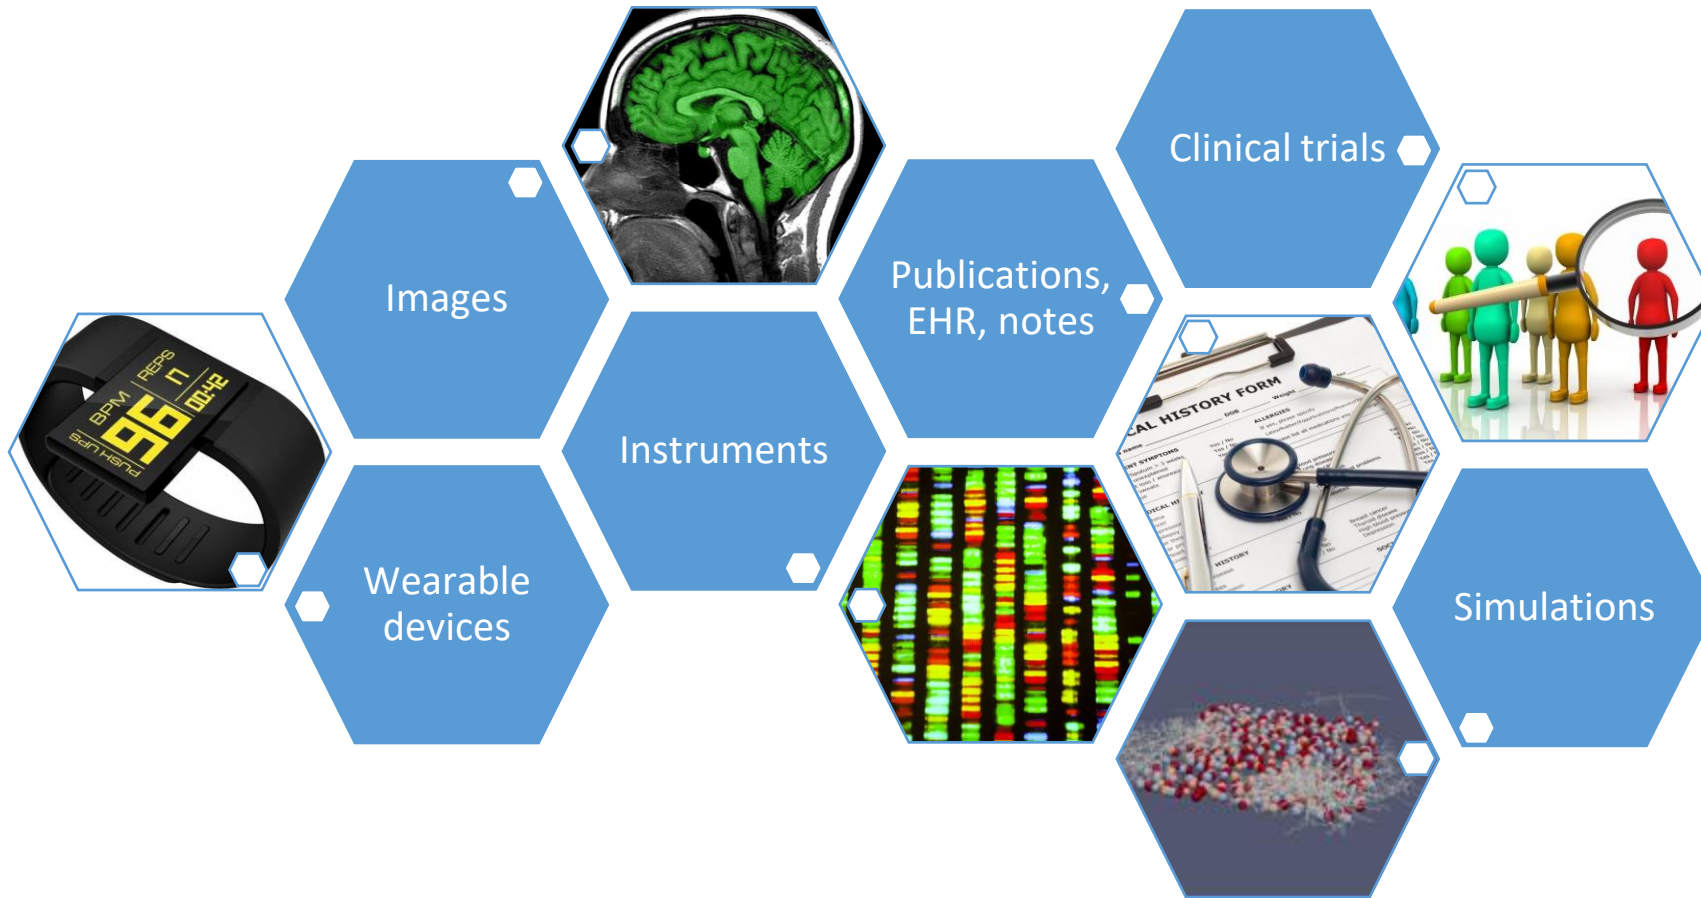
Further Information : Breitwieser et al. 2021 (<https://arxiv.org/abs/2006.06775>)

Extend the Agent-Based Engine

Agent-Based Social Simulation (ABSS)
Collaboration with statistics and demographics Institutes

<https://kt.cern/news/news/knowledge-sharing/cern-technology-support-study-socio-economic-inequities-new>

Computing for Society: Living Lab



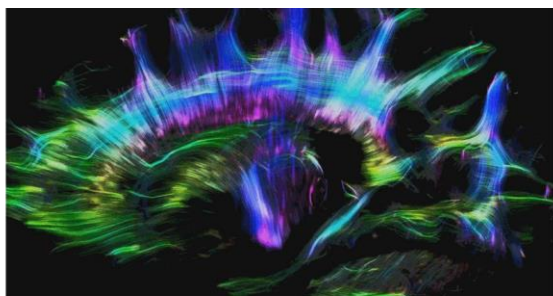
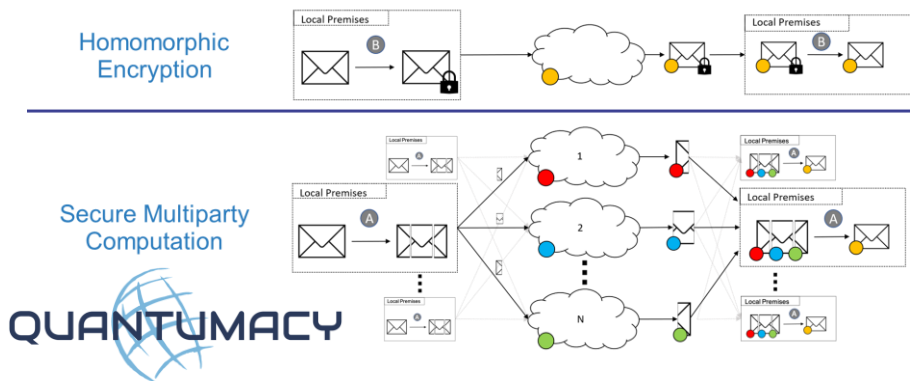
- “150 EBytes of medical data in the US, growing 48% annually” [1]
- Cost of instruments and laboratory equipment decreasing fast (e.g. **sub-1k\$ genomic sequencers**)
- Medical and fitness wearable devices on the rise, projected data produced in 2020 **335 PB/month** [2]

[1] Esteva A. et al., A Guide to Deep Learning in Healthcare, in Nature – Medicine, Vol. 25, Jan 2019, 24-29

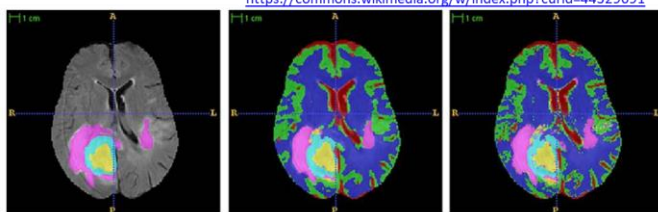
[2] <https://www.statista.com/statistics/292837/global-wearable-device-mobile-data-traffic/>

Data Analysis Applications to Medicine

Private Deep Learning for Healthcare



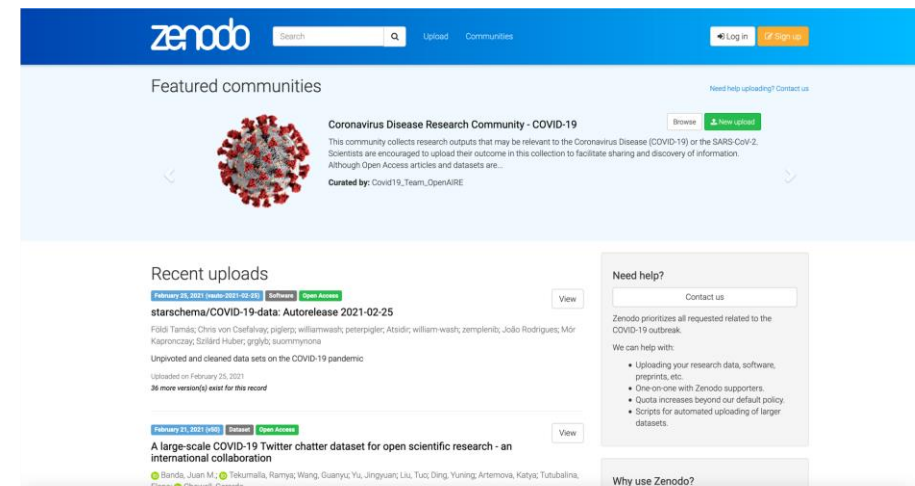
Source: [Alfred Anwander - https://www.youtube.com/watch?v=jrC8iY6_aZQ, CC BY-SA 4.0, <https://commons.wikimedia.org/w/index.php?curid=44329691>



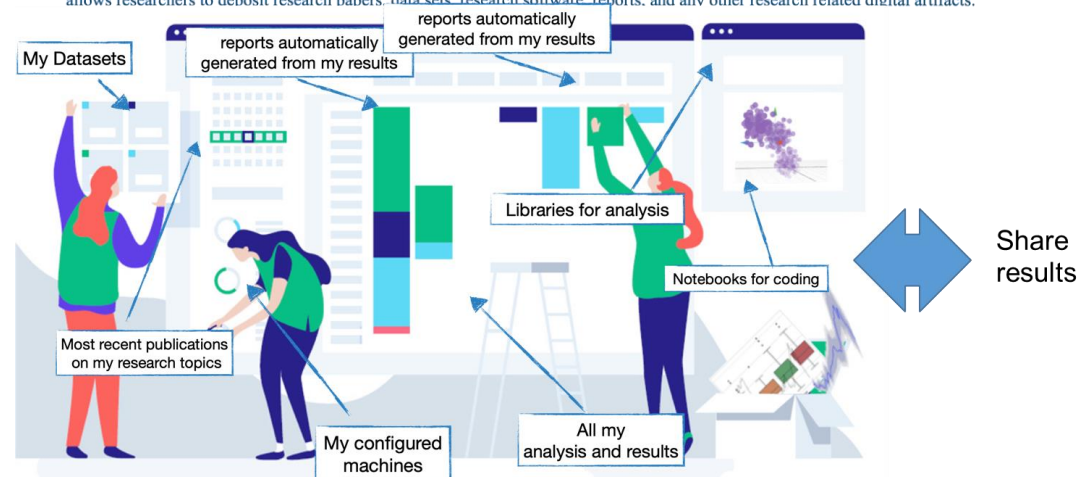
Source: [Bauer et al., 2011 MICCAI]

https://indico.cern.ch/event/1009424/contributions/424615/attachments/2205771/3731947/Private_Deep_Learning_for_Healthcare.pdf

CERN Science for Open Data (CS4OD)



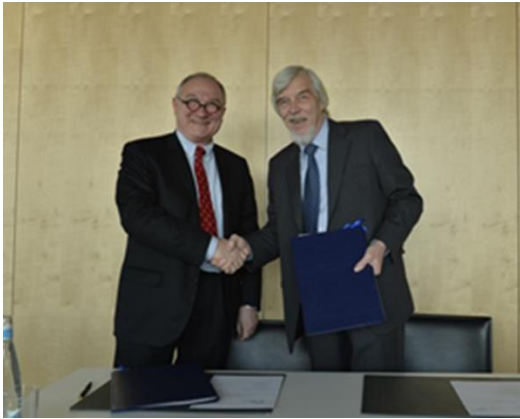
Zenodo is a general-purpose open-access repository developed under the European OpenAIRE program and operated by CERN. It allows researchers to deposit research papers, data sets, research software, reports, and any other research related digital artifacts.



CERN-ESA Cooperation Framework



“CERN and ESA have common roots and share a long history of pioneering research work in their respective fields. This cooperation agreement will foster synergies between the expertise, know-how and facilities available in the two Organizations.”



High level bilateral framework established since 2014 to foster synergies and explore collaborations in 12 different technology areas.

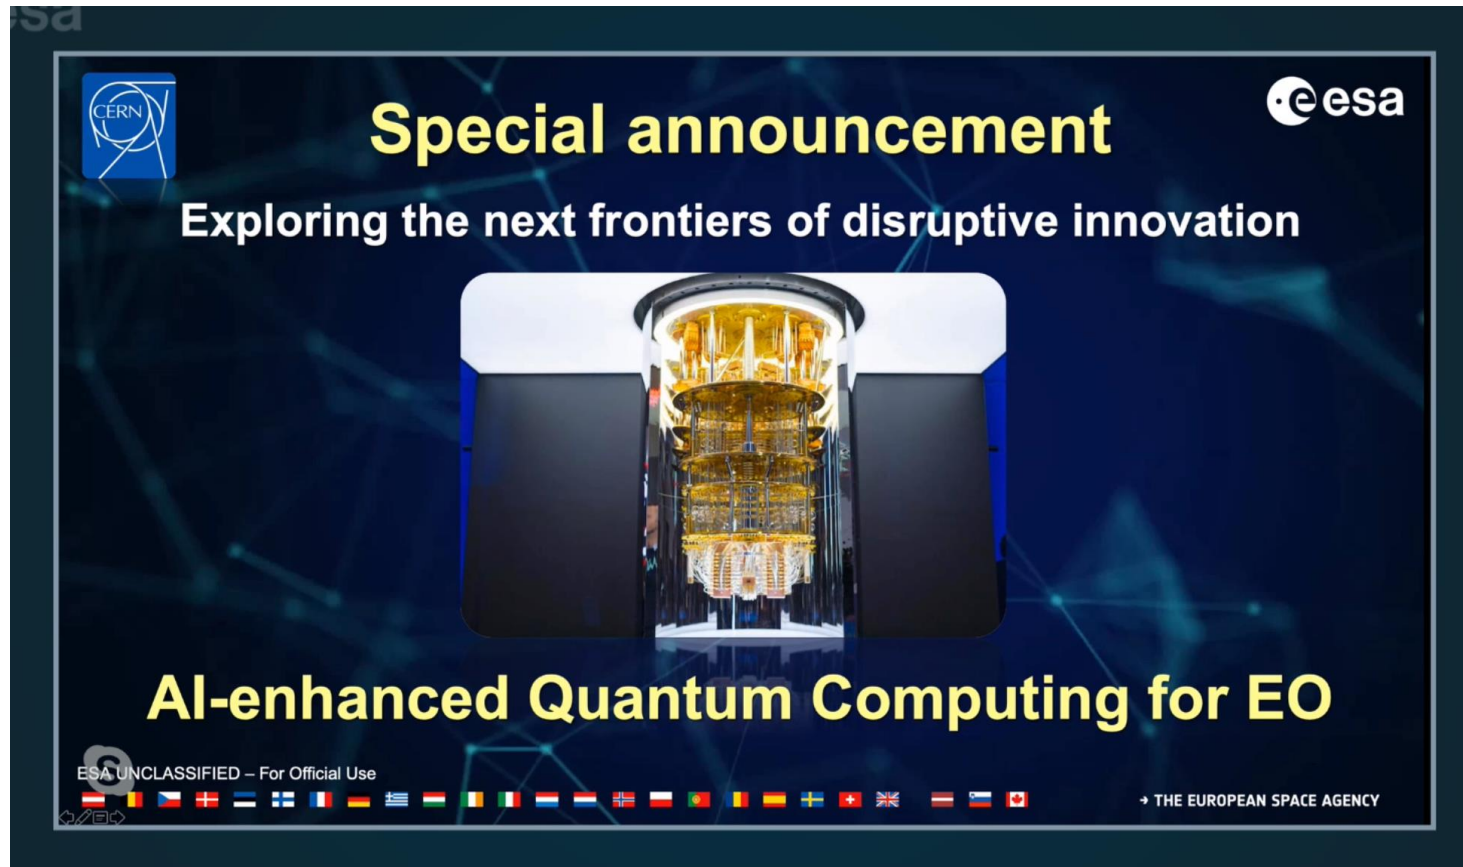
Implementing Protocol on «Radiation Environments, Technologies and Facilities » signed in 2019 including provisions for joint phds, facilitated access to testing facilities and 7 specific projects:

- 1 - High Energy Electrons Tests
- 2 - High Penetration Heavy Ions Tests
- 3 - COTS assessment strategy
- 4 - In Orbit Technology Demonstration
- 5 - Rad-Hard and Rad-Tol components and modules
- 6 - Radiation Detectors, Monitors and Dosimeters
- 7 - Simulation tools for radiation effects



Partnership on quantum computing to be implemented in 2021 for maximizing synergies in data mining and pattern recognition, and support EU Destination Earth initiative, aimed at creating an AI-driven dynamic, digital replica of our planet.

ESA Twin-Earth & QC4EO

A promotional graphic for the QC4EO project. It features a dark blue background with a glowing network of white lines. In the center is a photograph of a large, cylindrical, multi-tiered structure, likely a particle detector or quantum computing hardware, illuminated with warm yellow light. The text is in white and yellow. Logos for CERN and ESA are in the top corners. A row of European Union member state flags is at the bottom left, and the text 'THE EUROPEAN SPACE AGENCY' is at the bottom right.

Special announcement

Exploring the next frontiers of disruptive innovation

AI-enhanced Quantum Computing for EO

ESA UNCLASSIFIED – For Official Use

→ THE EUROPEAN SPACE AGENCY

QC4EO

Collaboration among CERN openlab, ESA Φ -lab and other research labs and universities

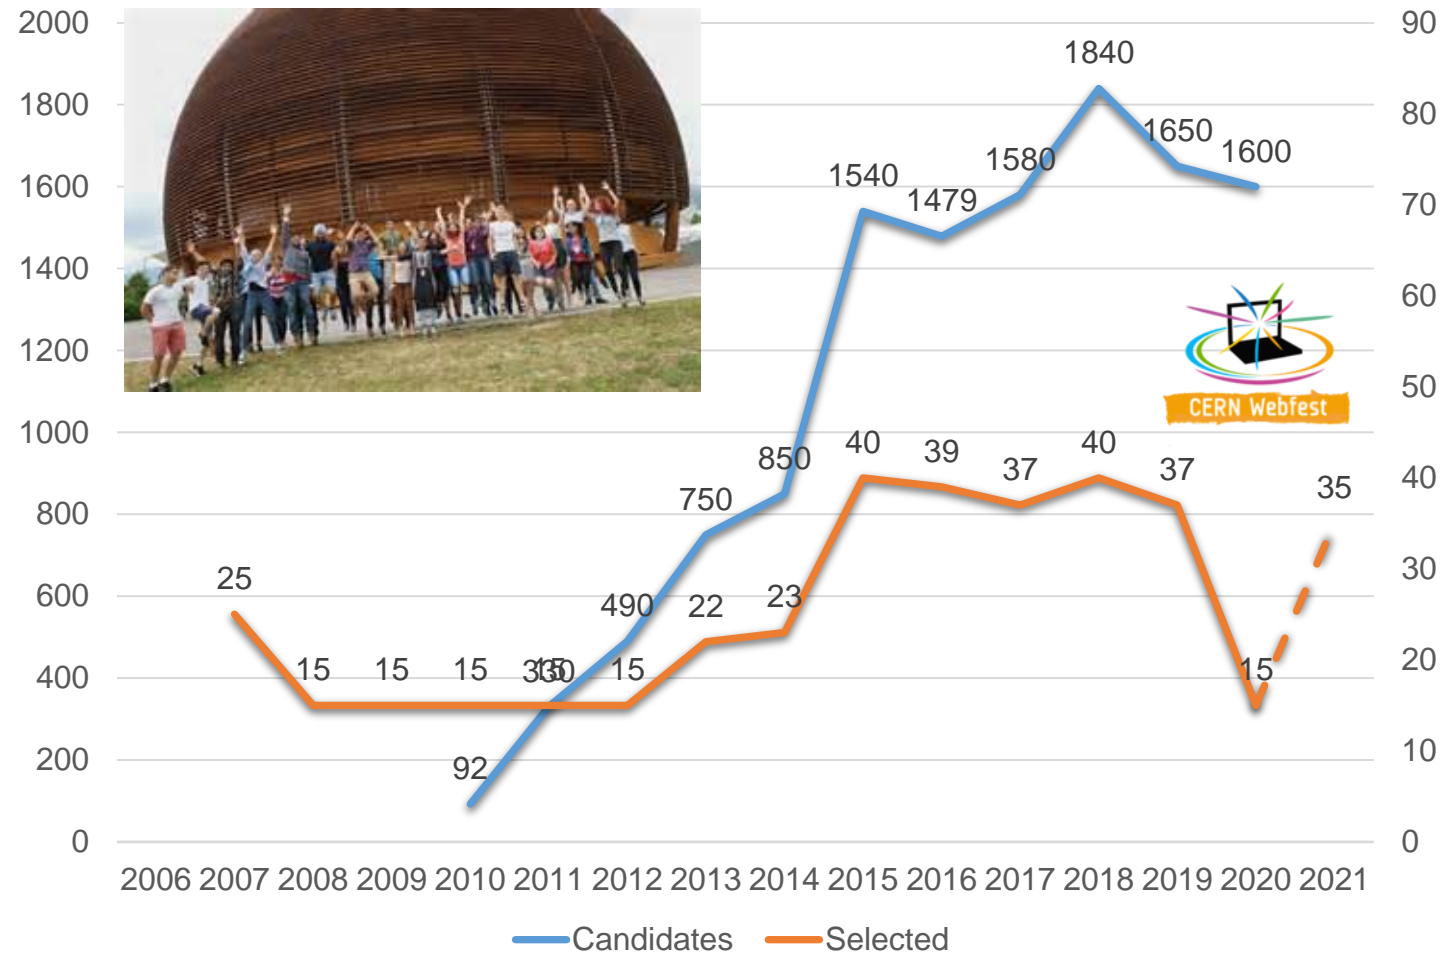
Investigation of impact of quantum computing and quantum machine learning at the intersection of Earth Observation and High-Energy Physics (image processing, data classification, error correction, etc.)

<https://phiweek.esa.int/>

Education and Training

CERN openlab Phase VII and Beyond

SUMMER STUDENT PROGRAMME



In 2019

- 1650 applicants
- 37 selected students
- Lightning talks session
- Technical reports

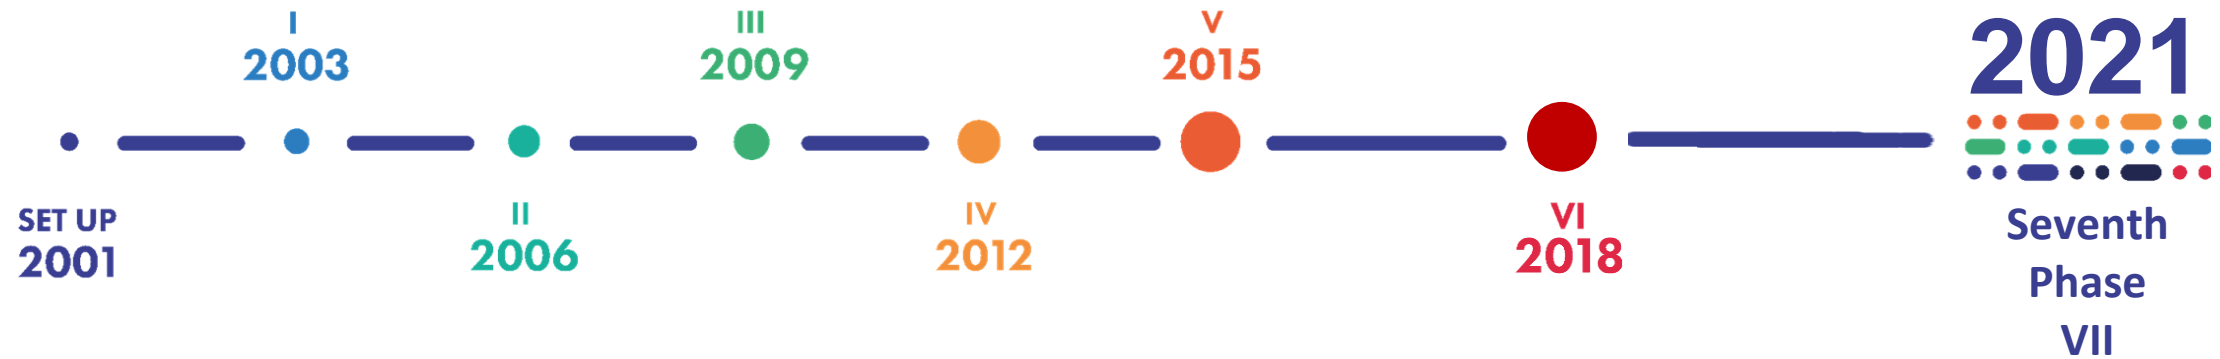
In 2020

- 1650 applicants
- 15 selected students
- Fully online
- Largest Webfest ever (>400 registered participants)

In 2021

- Grants for 35 selected students
- Fully online

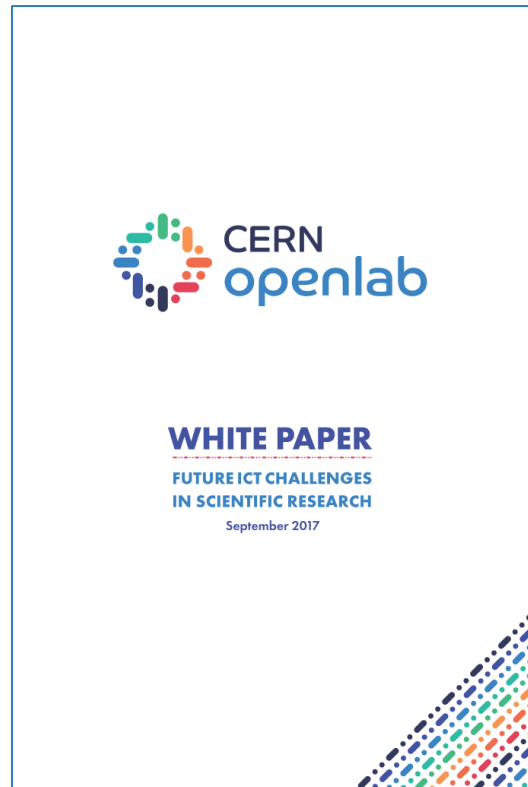
DRIVING INNOVATION **FOR 20 YEARS**



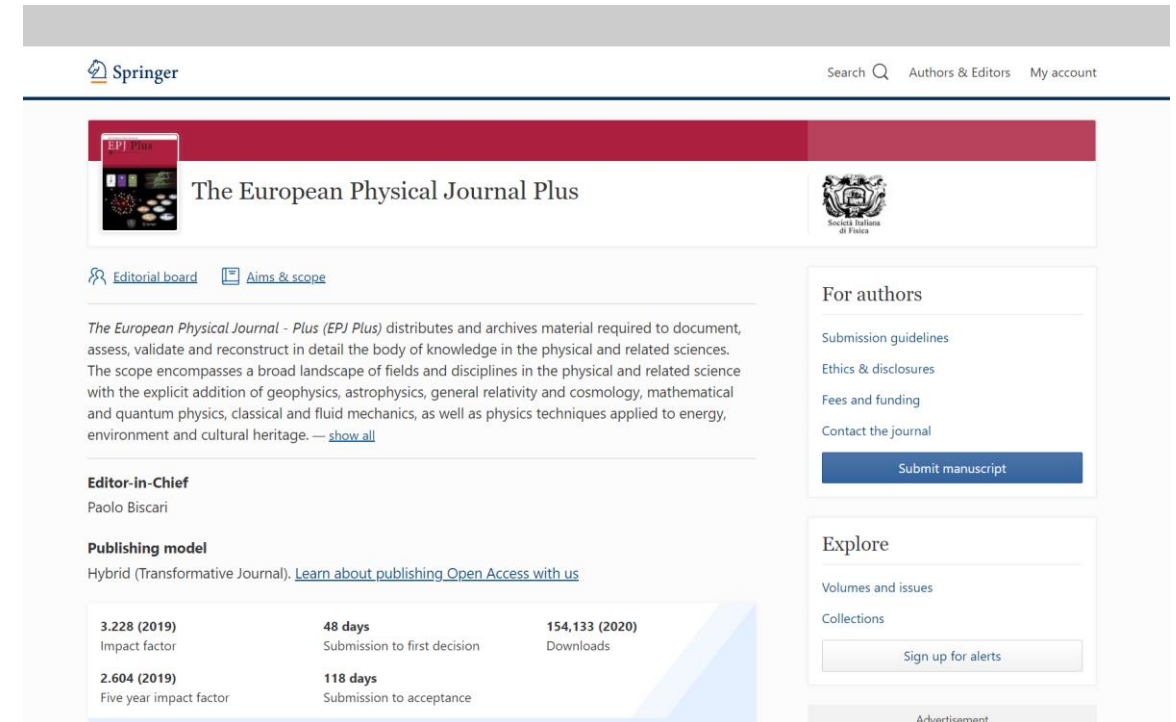
20th Anniversary Publication



2014 – Phase V



2017 – Phase VI



Phase VII Kick-Off event
in Q2-Q3 to present
more in details the
objectives and plans



Thanks!