



# Phase VII Overview

Alberto Di Meglio – CERN openlab Head

05/07/2021

# 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,  
1.2 B CHF budget

~3'200 staff, fellows, trainees, ...

>13'000 associates

1954: 12 Member States

**Members:** Austria, Belgium, Bulgaria, Czech republic, Denmark, Finland, France, Germany, Greece, Hungary, Israel, Italy, Netherlands, Norway, Poland, Portugal, Slovak Republic, Slovenia, Spain, Sweden, Switzerland, United Kingdom

**Candidate for membership:** Cyprus, Slovenia

**Associate members:** Croatia, India, Lithuania, Pakistan, Turkey, Ukraine

**Observers:** EC, Japan, JINR, Russia, UNESCO, United States of America


Numerous **non-member states with collaboration agreement**

>2'500 staff members, 645 fellows, 21 trainees

7'000 member states, 1'800 USA, 900 Russia, 270 Japan, ...

CERN

CMS




CMS

**1 PB/sec**  
of raw unfiltered data  
(> 2000 disks/sec)



LHCb

LHCb



ATLAS

ALICE



ALICE


CERN opens Phase VII and Beyond

ATLAS

# The CERN Data Centre in Numbers

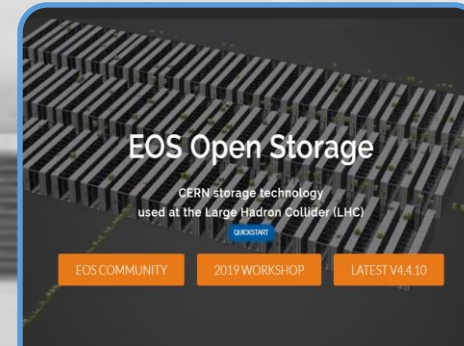



10 000  
Servers




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

375 000  
Cores



280 PB Hot  
Storage




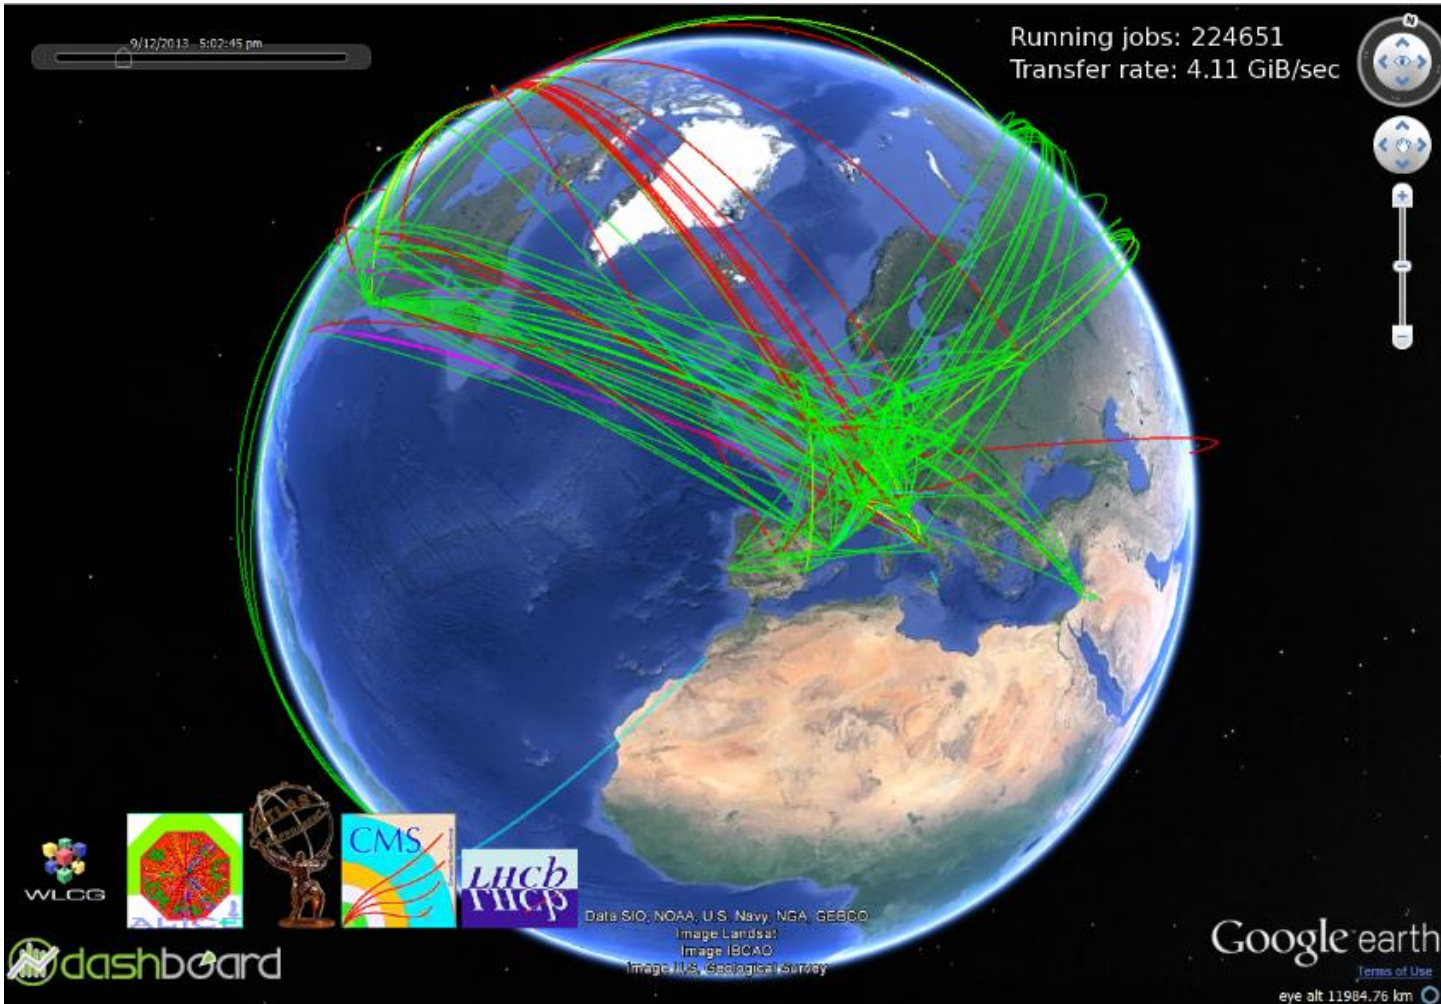
380 PB Cold  
Storage



35 000 km  
Fiber Optics



# Worldwide LHC Computing Grid



Tier-0 (CERN):

- Data recording
- Initial data reconstruction
- Data distribution

Tier-1 (13 centres):


- Permanent storage
- Re-processing
- Analysis

Tier-2 (42 Countries, ~170 centres):

- Simulation
- End-user analysis

- ~800,000 cores
- ~800 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

SWITZERLAND

FRANCE

FCC

100 KM LONG




# 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.

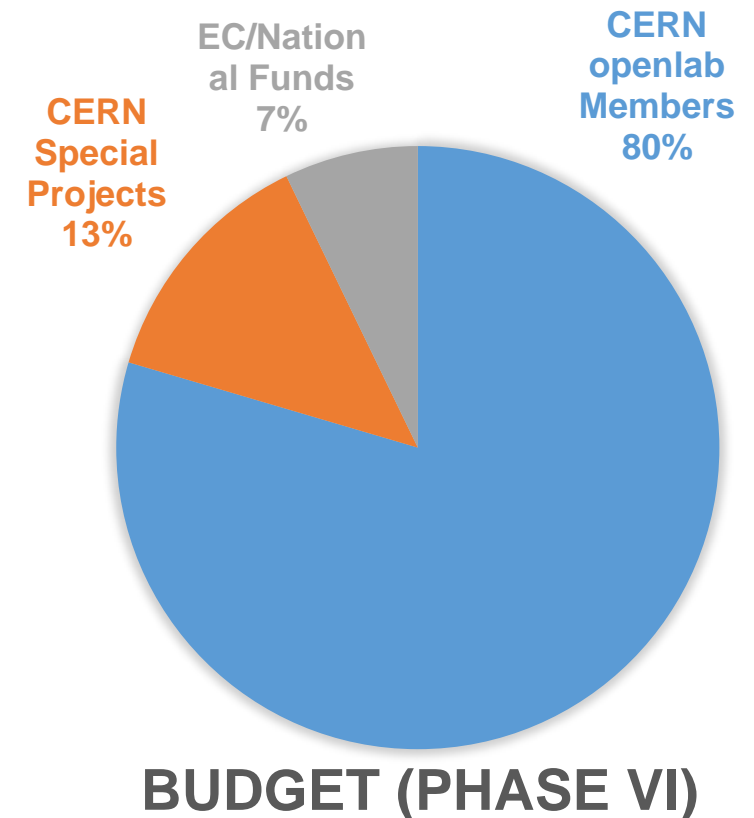
# DRIVING INNOVATION FOR 20 YEARS



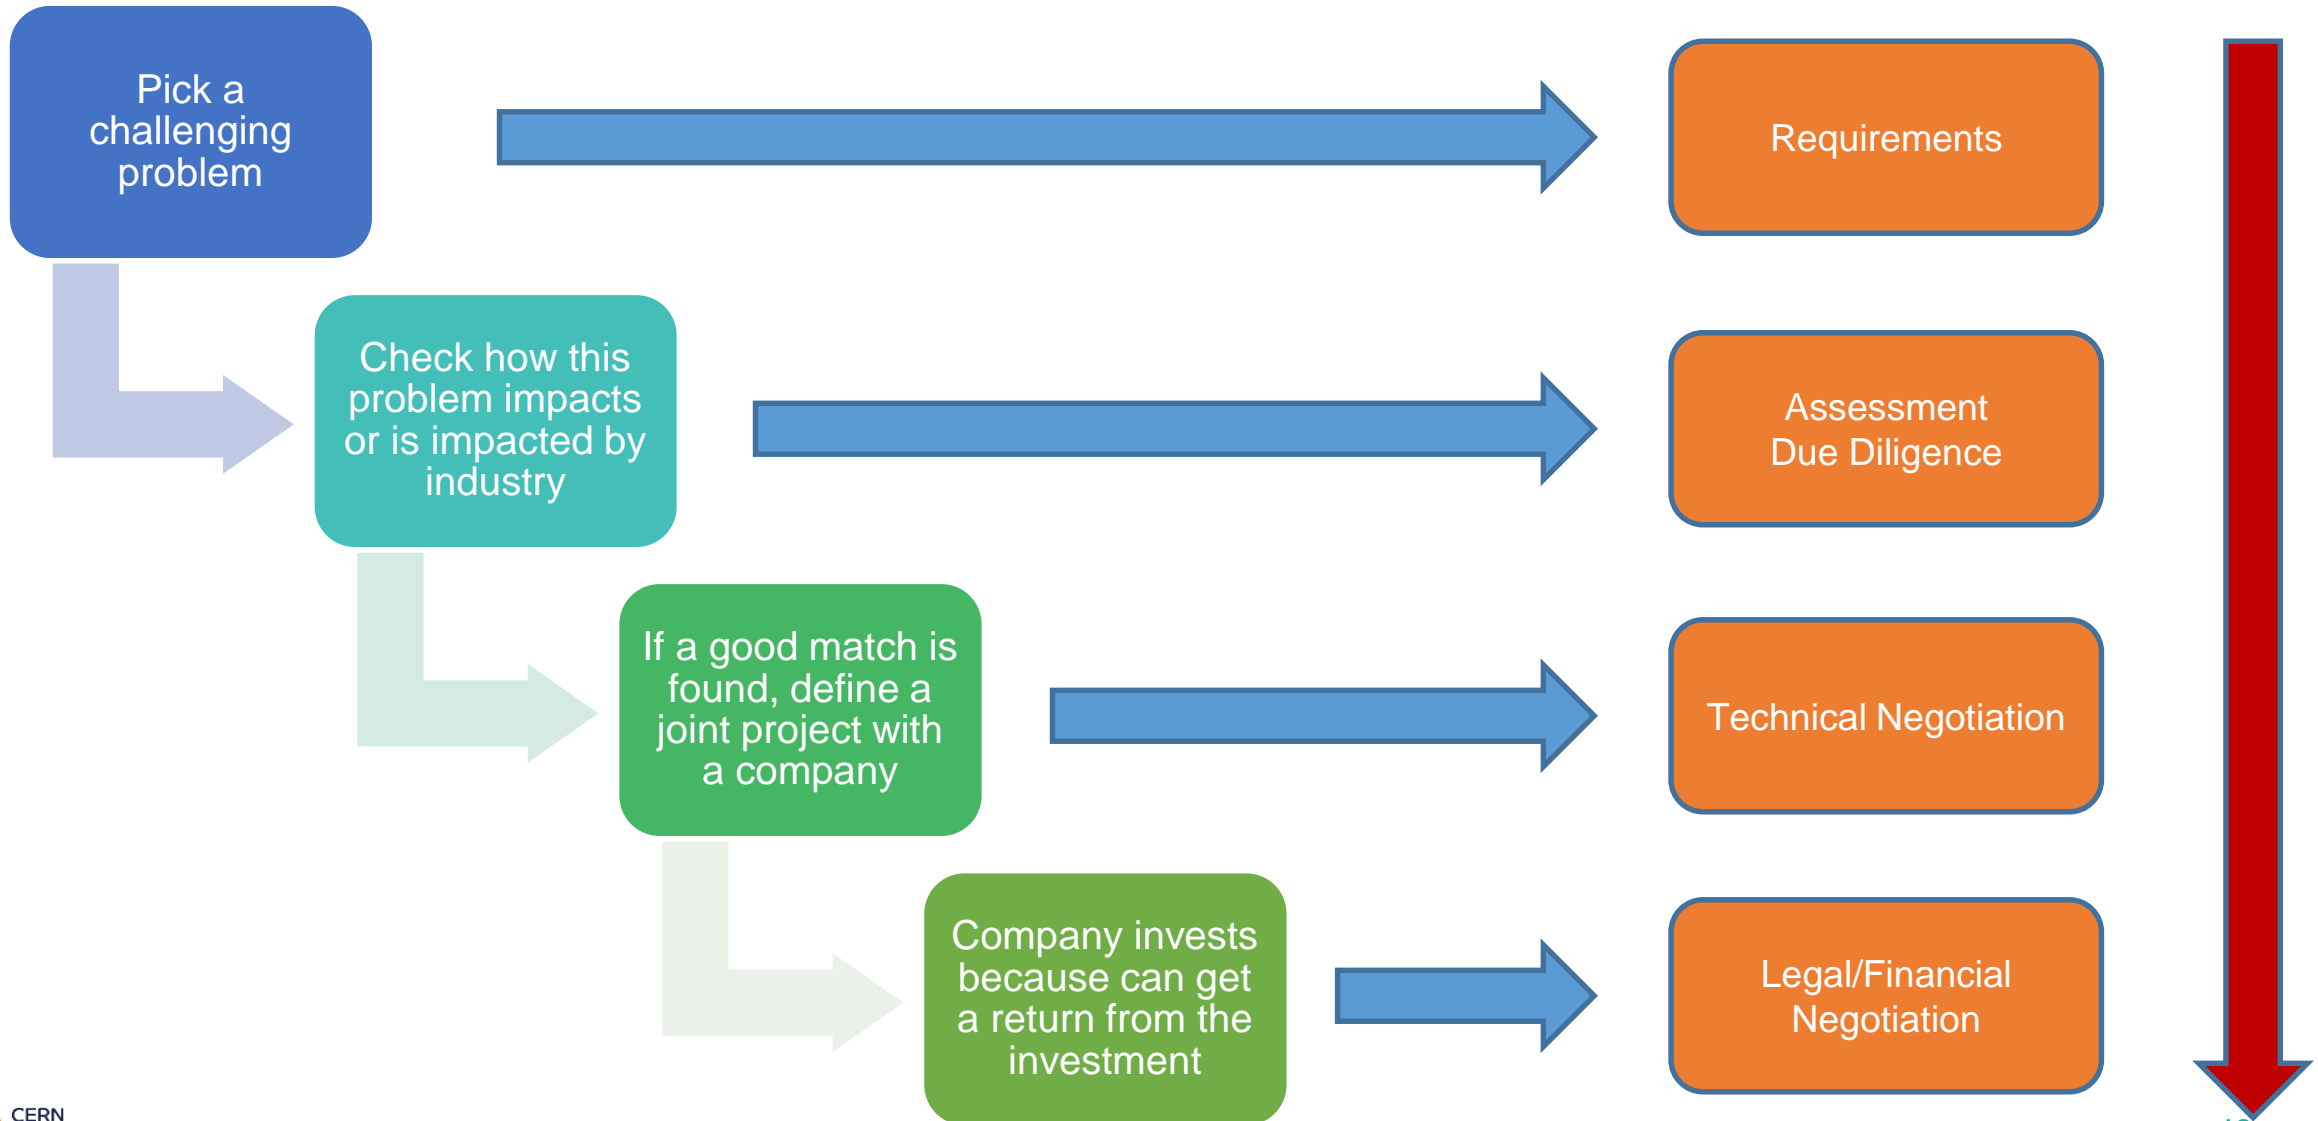
# CERN openlab in Numbers (Phase VI)

*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 across IT, EP, BE
- > 100k visitors to website
- > 5k applications to summer-student programme
- ~ 300 press articles
- > 100 main communications (by CERN and members)
- ~10 peer-reviewed papers per year
- Over 17k followers on our main social media channels.
- CERN openlab channel in the CERN's alumni platform



# Collaboration Model




# 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



HPC, Cloud, Quantum





Predictive/proactive maintenance and operations



Cloud technology, containers, scalability



Software Defined Networks, Security



Storage architectures, scalability, monitoring




CERN openlab Phase VII and Beyond



# Phase VI: Three Main Areas of R&D




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 Projects and Results

## Data Centre Technologies and Performance

Evaluation of heterogeneous architectures (Intel, IBM, E4 on NVIDIA architectures, Micron, EU-funded project DEEP-EST)  
 Unified programming models (investigations with Intel OneAPI)  
 Performance evaluation and benchmarking  
 Cloud IaaS/PaaS solutions with various providers (Oracle, Google, Microsoft)  
 Storage solutions

## Scale Out Capacity

ML/DL training and data analysis at scale with Kubernetes on Google Clouds  
 Large-scale 3D-GAN training/inference with Intel architectures at SURFsara  
 Satellite image processing, segmentation and generation with UNOSAT on T-Systems OTC  
 Initial assessments of HPCaaS, MLaaS and QCaaS solutions on multiple clouds in collaboration with IT-DI-EFP

<https://indico.cern.ch/event/1009424/> - 34 talks, 281 registered attendees, peaks of 130 connected

## New Computing Paradigms

ML/DL on classic and non-classic (non-Von-Neumann) architectures  
 Efficient use of new hardware architectures and algorithms  
 New algorithms for neutrino experiments with Dune and IBM  
 First Quantum Machine Learning pilots for HEP applications  
 Spin-off of Quantum Computing into the new Quantum Technology Initiative

## Working with Other Sciences

Expansion of BioDynamo from local project to open-source collaboration, more use cases (cell dynamics, cancer modelling, in-vitro simulation, epidemiology)  
 Extension of ML/DL algorithms for privacy (HE, SMCP)  
 Open Data Collaborations for Healthcare and Social Sciences  
 Earth Observation applications with UNOSAT and ESA

# Phase VII Programmes

## XT eXascale Technologies

A comprehensive investigation of HPC and Cloud infrastructures, platforms, 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<sup>2</sup> Quantum Technology Initiative – Computing & Communications

Assess the potential impact of quantum computing for science, investigate QML algorithms and areas of potential quantum advantage, collaborative quantum computing (simulation) platform

## MSC Multi-Science Collaborations

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



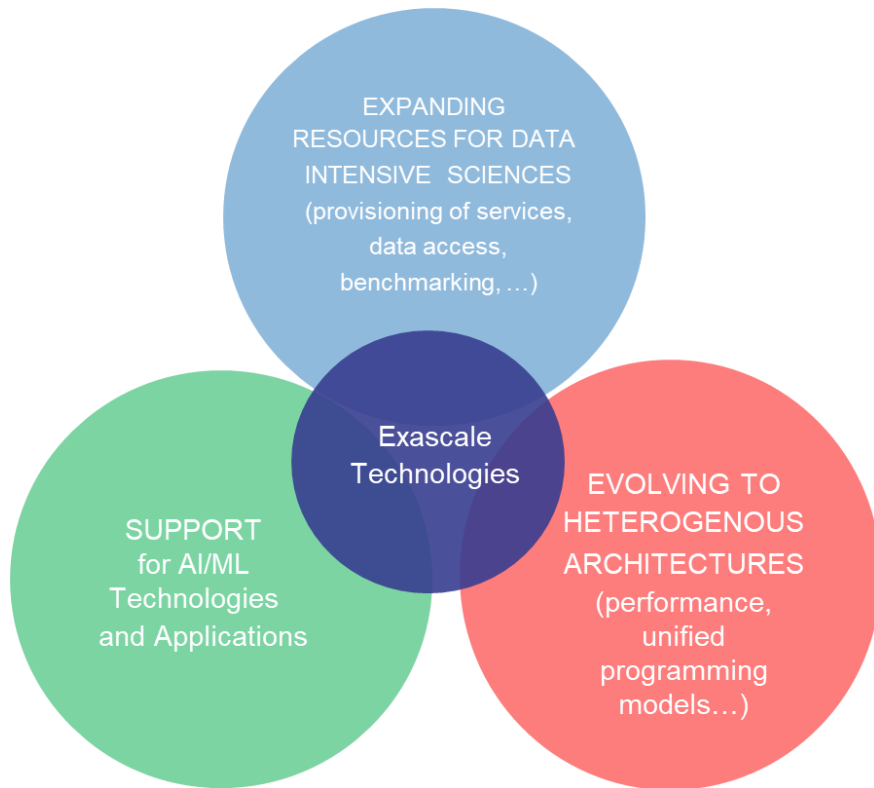
# Phase VII: New Challenges

Phase VII coincides with the LHC Run 3, it's an opportunity to consolidate the investigations started in Phase VI and move 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.  
*How to develop AI applications in HEP research? For which activities? How to deploy and support them?*
- 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 for international scientific communities?*

# eXascale Technologies (XT): HPC

*Lead: Maria Girone*



## Key Collaborators

*HPC Collaboration* CERN, SKAO, GÉANT and PRACE

Benchmarking on HPC

Data Access on HPC

AAI on HPC

*CERN openlab industry members*

**Intel** OneAPI, DAOS

**Micron** DLA

**IBM** Power

**E4** Innovative Solutions with Nvidia in HPC and AI

*Collaboration with WLCG, LHC experiments, EP-SFT, EP-DT*

## Main Projects/Activities

*EU funded projects*

- **CoE RAISE** (Research on AI- and Simulation-based Engineering at Exascale)
- **EGI-ACE: Advanced Computing for EOSC** - HTC to HPC transition

*Collaborations (ongoing and in discussion) with Universities and Research Centers*

- EU: NTNU HPC-lab, SURFSara, LRZ, CINECA, BSC
- US: Flatiron Institute, University California San Diego, University of Iowa

# eXascale Technologies (XT): Clouds

*Collaboration with IT-DI-EFP*



## Key Collaborators

*IaaS/PaaS providers: Oracle, Google, Microsoft, Amazon, T-Systems*

*Through EU CloudBank (preferred mechanism) or direct agreements*

## Main Projects/Activities

*Deploy of a series of use-cases in the domain of Artificial Intelligence, Machine Learning and Exp./Th. Physics Simulation, based on:*

- *Benefit to the scientific programme of the Large Hadron Collider and LHC Experiments*
- *Need of resources that are not available or not cost-effective on premise*

*Collect and share information on impact, feasibility, sustainability, cost for future production support*

# Artificial Intelligence for Science (AI-S)

Lead: Sofia Vallecorsa

## Key Collaborators

Intel, Google, IBM, Microsoft, SURFsara, LRZ, BSC

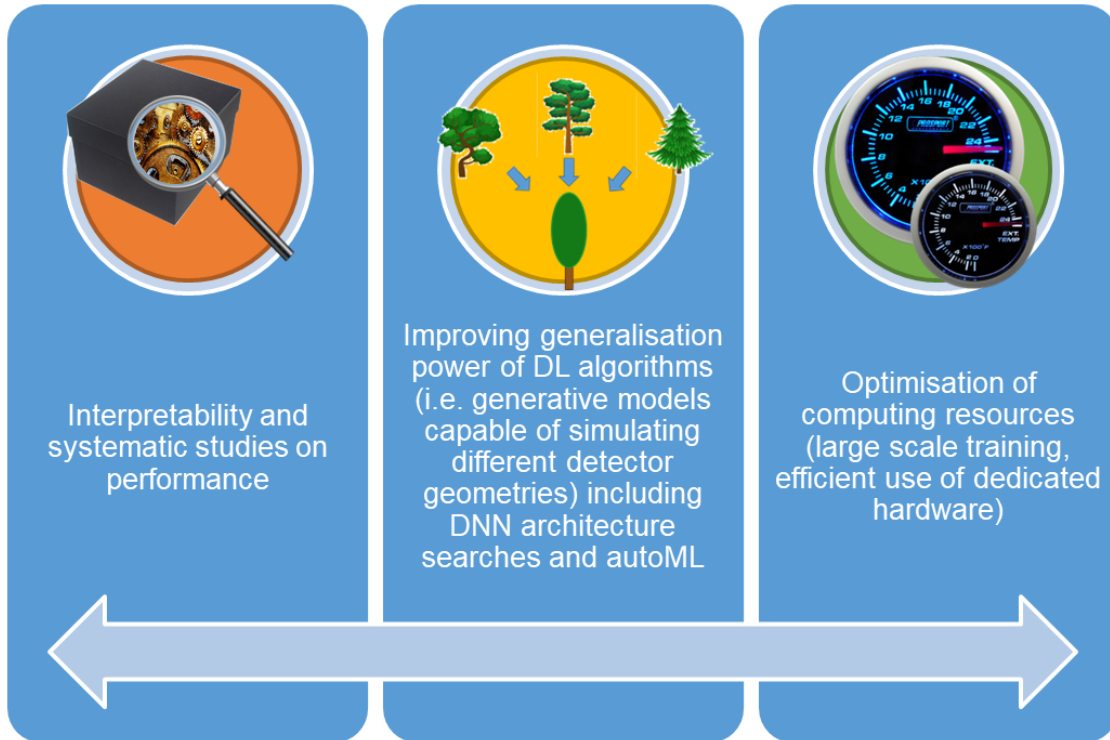
## Main Projects/Activities

Projects (ongoing or in discussion)

- *Boosted 3D-GAN, Low-Precision representation for DL for fast detector simulation (Intel)*
- *Anomaly Detection for physics and engineering (rare events, accelerators)*
- *Deep-Learning-based optimization of network traffic (with IT-CS)*
- *Parallel training of DL algorithms with Google TPUs*
- *ML/DL and Monte Carlo models acceleration with Graphcore IPU on MS Azure*

Areas of R&D

- **Deep Learning models development** for different steps of the data processing and analysis chain: simulation, anomaly detection, reconstruction and real-time data processing.
- **Computing resources optimisation for DL: highly parallel architectures** (TPUs, IPU, ...) and HPC clusters multi-GPU and multi-CPU
- **Investigations about the DL models generalisation capabilities and the systematic uncertainties** related to the training process and training samples.
- **Auto-ML studies:** study automatised procedures to design the best DL architecture for the specific dataset



# Quantum Computing and Communication (QTI-C<sup>2</sup>)

Leads: Sofia Vallecorsa (Computing), Alberto Di Meglio (QTI Coordination/Communication)

## Key Collaborators

CQC, IBM, gluoNNet, METU, ICEPP, DESY, SNUBH, IDQ, ETHZ, U. Oviedo, U. Waterloo, CERN-BE

## Main Projects/Activities

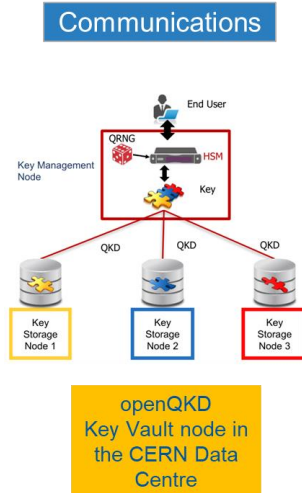
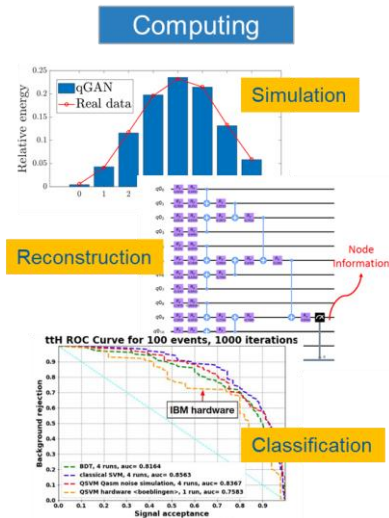
Quantum Communication and Computing projects are developed in the context of the CERN QTI

## Projects

- Quantum classifiers/anomaly detection models (DESY, ETHZ, ICEPP)
- CMS Patatrack Qtrack (Waterloo)
- qGAN models (EPFL, ESA)
- qGNN for tracking (ATLAS, gluoNNet, METU)
- Beams steering optimization (with BE, D-Wave, MS)
- QKD and time synchronization for quantum networks (openQKD, White Rabbit)
- Design and deployment of the Quantum Computing (Simulation) Platform (IBM, Atos, E4, INFN, Oviedo, CINECA, plan to grow a distributed collaboration), transition to “production-like” over the next three years

## Main areas of R&D

- Development and characterisation of QML models, quantum simulations, Quantum Reinforcement Learning, algorithms for clustering, noise and error mitigation techniques, quantum computing platforms



# Multi-Science Collaborations (MSC)

Leads: Fons Rademakers (BioDynaMo), Alberto Di Meglio (Living Lab, Earth Observation), Sofia Vallecorsa (Earth Observation)



Knowledge Sharing



Joint Research



Industrial/Social Impact

## SUSTAINABLE DEVELOPMENT GOALS



## Key Collaborators

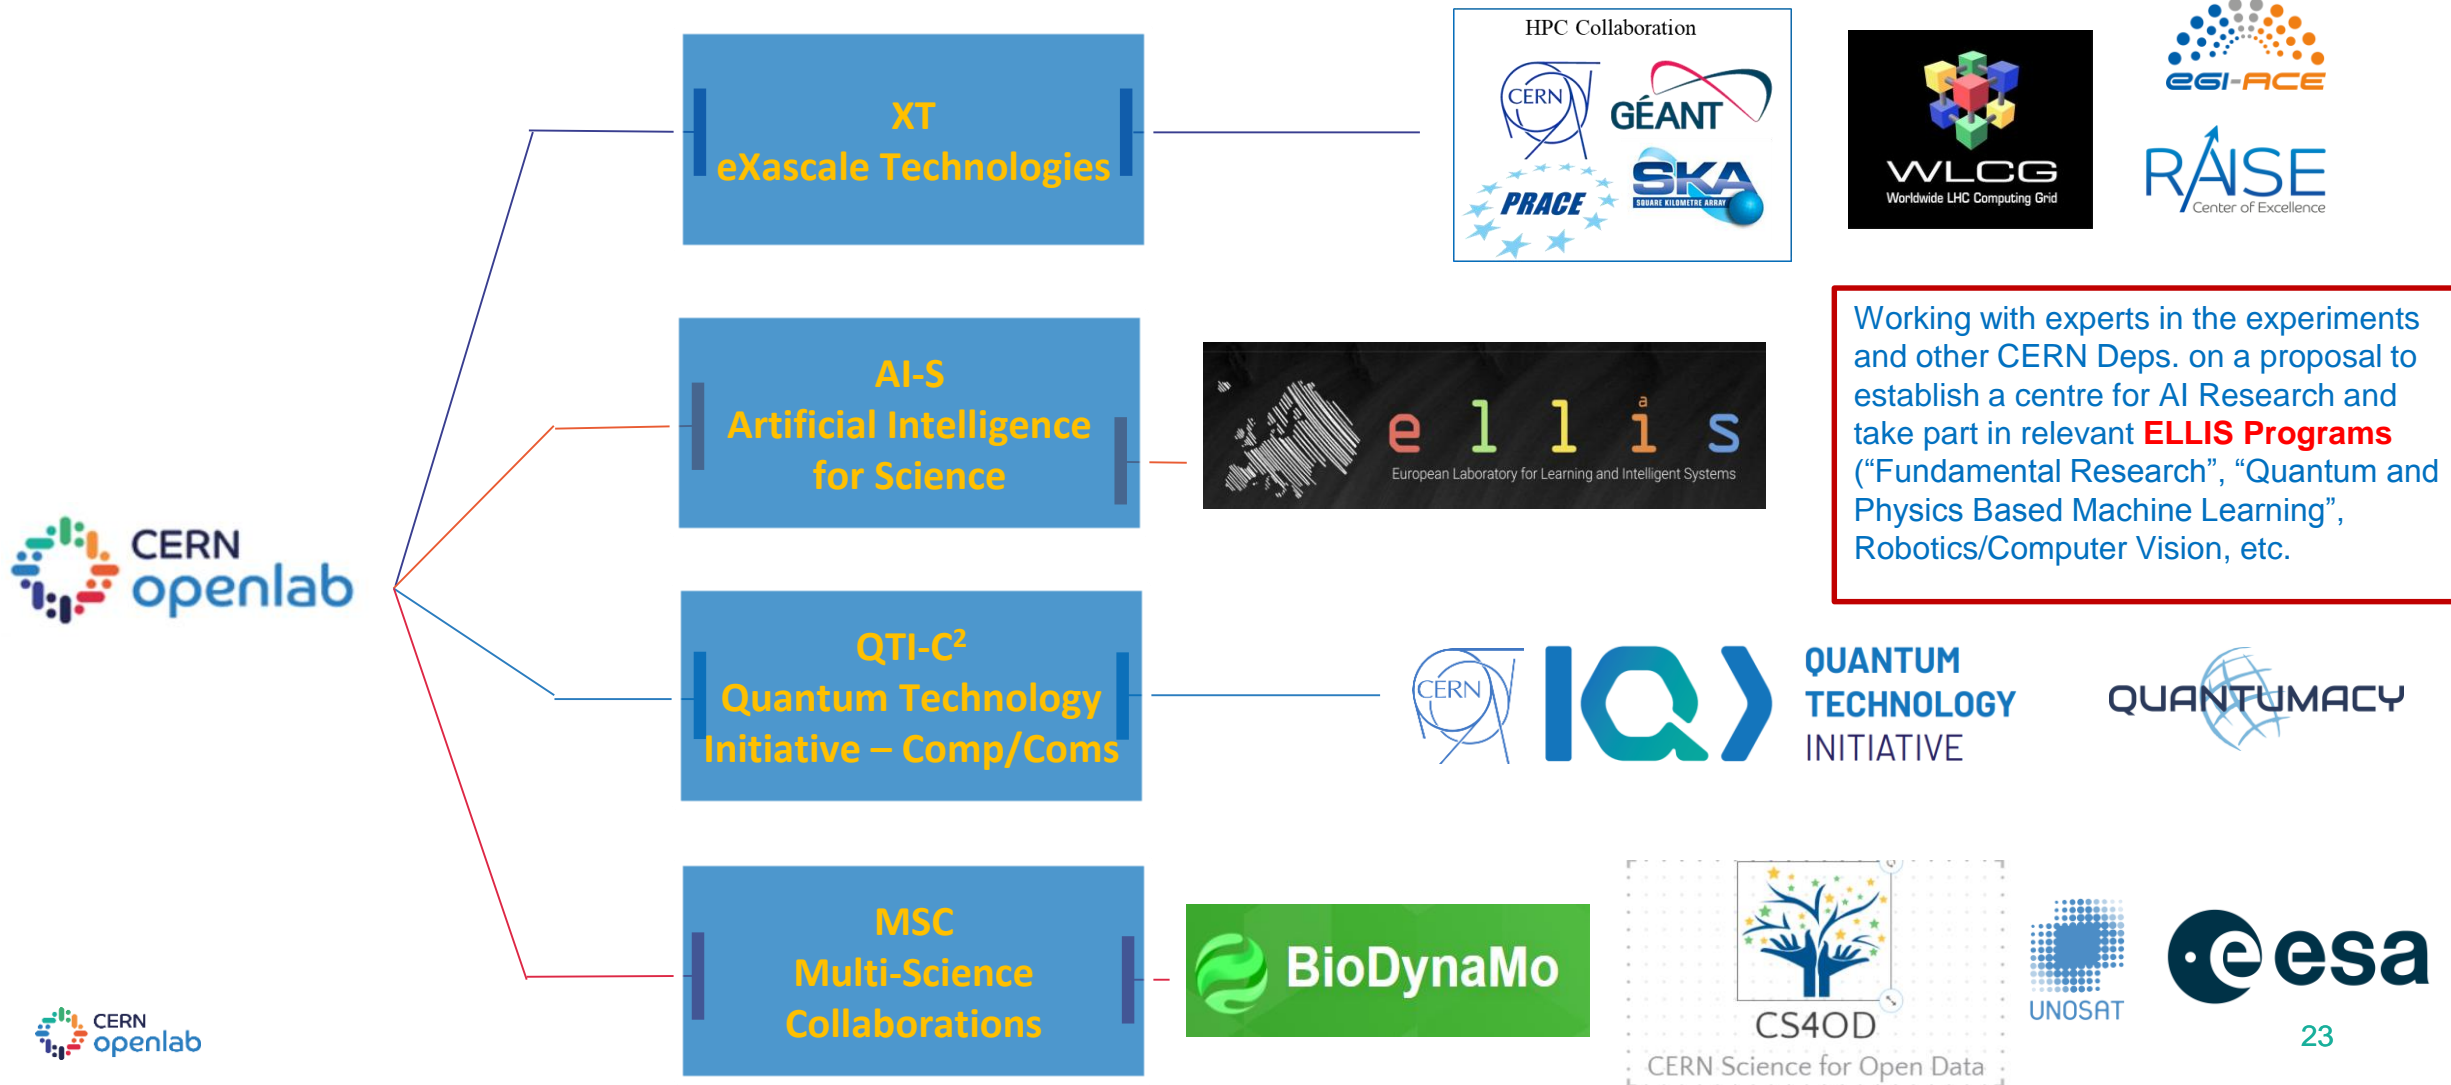
ESA, LRZ, TUM, DLR, ECMWF, UNOSAT, Uni Florida, Bocconi, Milano Bicocca, Uni Madrid, SNUBH

## Main Projects/Activities

- **BioDynaMo**, agent-based simulation for medical and social sciences research
- **Living Lab**, health/privacy applications of ML/DL approaches (Parkinson's detection, large-scale confidential computing)
- **CS4OD**: Data fusion/analysis over open-source data for medical/social/demographic applications
- **Quantumacy**, QKD infrastructures
- **ESA**: Development of QC-based algorithms for Earth Observation
- **UNOSAT**: Deep Learning for EO, development and optimisation of models targeted to landscape recognition, seismic building damage assessment, etc. from satellite imagery - Deep Generative models for the creation of synthetic satellite image datasets to improve training

# Framework Initiatives and Projects

Frame the CERN openlab programs within broad initiatives or projects at CERN, HEP or international level to ensure alignment, impact and adoption



# Computing/Storage Resources

## Collaboration with IT-CF

### On-premise



Intel Xeon nodes with different processors  
Ice Lake coming in Q2

E4/Nvidia nodes with T4 and V100



IBM Power 8 and 9 clusters

Intel Optane equipped nodes



Classic general-purpose or optimized nodes for quantum simulation

## Collaboration with IT-DI-EFP/IT-DB/IT-CM

### IaaS/PaaS MLaaS/HPCaaS/QCaaS



### HPC

Access to PRACE Tier0 testbed facilities (through HPC Collaboration)

San Diego Supercomputer Center (SDSC)  
CoreSite NY2

Iceland Data Center AtNorth  
SurfSara, CSCS, Cineca, LRZ



Norwegian University of Science and Technology





# Communications and Outreach

*Lead: Andrew Purcell*

- Communications support for 34 joint R&D projects
  - Projects distributed across departments and experiments.
  - Liaison with 20 external companies/institutions on all comms matters.
- Promotion of education/training activities
  - Summer-student programme.
  - Webfest (will be public, online event again this year).
  - Public talks, workshops, etc.
- Promotion of the CERN openlab concept
  - Further increasing awareness of mechanism/opportunities within CERN, across organisation.
  - External promotion: value of collaboration, knowledge sharing, new opportunities for collaboration, areas of mutual interest, etc.
  - All work summarised in annual report (next version due later this quarter).
  - Special activities to mark CERN openlab's 20th anniversary later this year (journal issue, event, and more).


# Education & Training

*Shared responsibility - Miguel Marquina directly responsible for integration with existing CERN programmes*

- 3 Summer Student Programmes (2 in presence, 1 online)
  - Lecture programmes and visits
- 16 Workshops and training events
  - ML/DL, CUDA, oneAPI, Quantum Computing, Computing for healthcare
- 2 Stages d'observation for high-school students (1 in presence, 1 online)
- Participation in the CERN Diversity and Special Needs internship and technical student programmes

# Summer Student Programme

Lead: Kristina Gunne



## 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 ~30 selected students
- Fully online



 **Thanks!**

*alberto.di.meglio@cern.ch*

@AlbertoDiMeglio