CERN openlab

CHEP 2024

A technical overview of industry-science R&D projects for the High Luminosity LHC under CERN openlab

Thomas James (CERN)

on behalf of the CERN openlab technical team

*Maria Girone (head), Luca Atzori, **Thomas James**, Antonio Nappi, Luca Mascetti*

CERN OPENLAB MISSION

Four primary missions

1
Establishing strategic industry collaborations

2
Fuelling technological innovation

3
Exposing technology to researchers

4
Nurturing knowledge and growth in young researchers

Current technical and communications team

Primary role

To act as conduit and facilitator for collaboration in computing, science, and technology between:

CERN openlab

The science communities (CERN departments and groups; R&D teams at CERN; research centres) ⟷ **Technology Providers** (industry)

HL-LHC CHALLENGES

- Current data analysis ecosystem under-equipped to handle the expected increase in data generation and complexity of High-Luminosity LHC.
- Based on current technological evolution, offline processing will fall an order of magnitude short of the demand of HL-LHC.
- Must embrace new hardware architectures, and heterogeneous hardware infrastructure.
- Must make significant investment in testing and integration of novel hardware architectures produced by both the research community and industry.

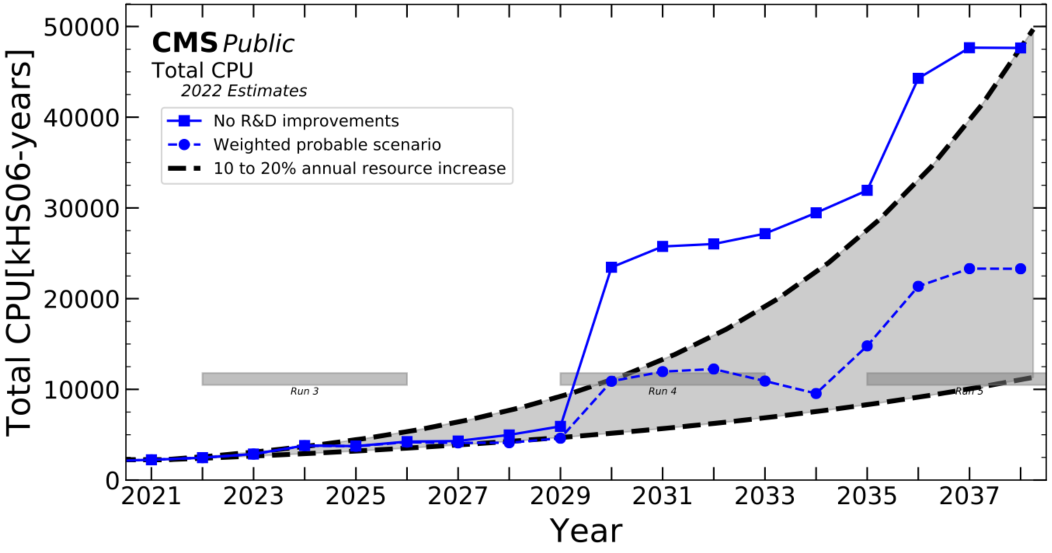

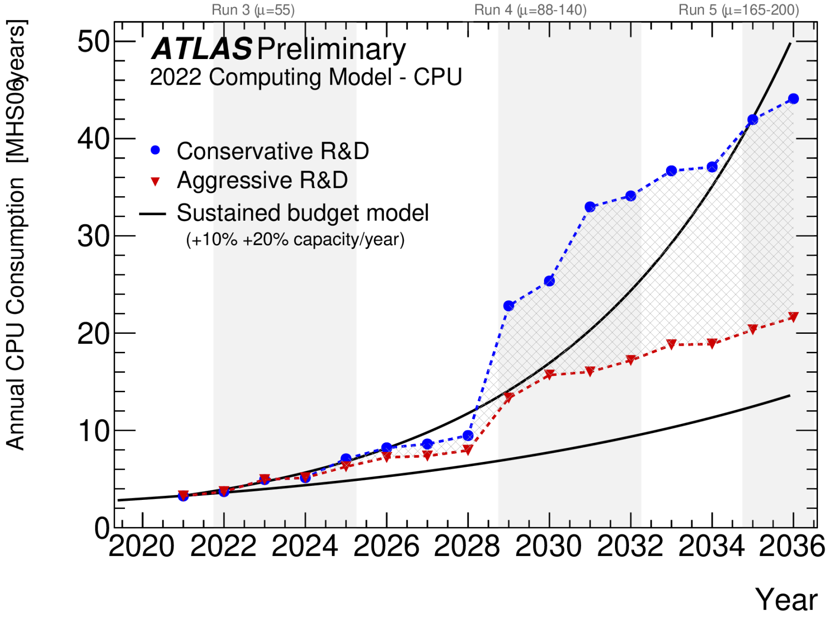

CERN openlab

CERN OPENLAB PHASE VIII

Structured three-year phase cycles:

- systematically assess technological evolution
- anticipate future needs
- delineate overarching thematic priorities.

Establishing a managed portfolio of small to medium-size, agile projects with technology providers with clear impact on the CERN IT Technology Roadmap.

Identifying a few collaborations, especially at the level of the computing infrastructures, of high potential impact and act as an initial incubation step for longer-term collaborations.

OBJECTIVES

High-level: Accelerating Computing for Science

Pioneering sustainable and emerging computing and storage solutions

Harnessing heterogeneous computing and AI for a greener future

Fostering synergies and technology transfers between industry and sciences

Sustainable Infrastructures

- Heterogeneous computing platforms and infrastructures
- Computer architectures and software engineering
- Storage and data management
- Artificial intelligence algorithms, platforms and applications
- Applications for society and environment

Emerging Technologies

- New materials for long term digital storage
- Digital twins
- Quantum computing and networks

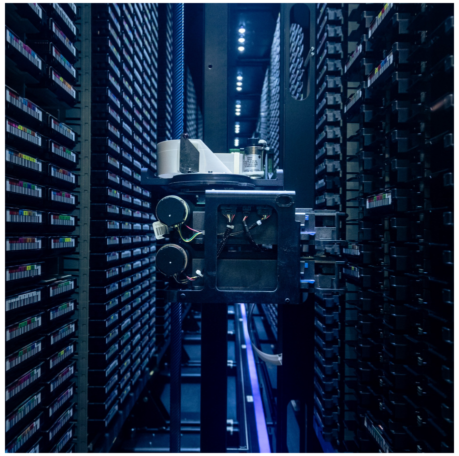

CERN openlab

CURRENT INDUSTRY AND RESEARCH MEMBERS

SIEMENS, micron, E4 Computer Engineering, Oracle, cerabyte, Università degli Studi di Trieste, Simons Foundation, INFN, intel

PRE-AGREEMENT STAGE

Pure Storage, Pasqal, NVIDIA, Johnson & Johnson

In discussions and negotiations with ~10 more

HETEROGENEOUS ARCHITECTURES

In collaboration with

intel NVIDIA E4 COMPUTER ENGINEERING

Hardware landscape becoming increasingly heterogeneous
- varied, often more specialised and high-performing architectures entering market.

CERN openlab has established a heterogeneous architectures testbed
- provides a rich ecosystem to access and evaluate novel architectures
- on-premise and remote resources.

Heterogeneous Architecture Testbed: Mandate

- Reception, installation, configuration, benchmarking of new hardware
- Hardware and software maintenance
- Project and user support: access grant and system administration
- Support for tools on CVMFS
- Organisation and support for workshops, training courses

HEPscore benchmark allows for performance comparison across architectures
- reported to, and analysed by, industry and research partners.

Heterogeneous Architecture Testbed: Hardware

- CPU
  - Intel Xeon Gen 6, 5, 4 …
    - *Sierra Forest*
    - *Emerald & Sapphire Rapids*
  - IBM Power 8 & 9
  - AMD EPYC*: Genoa, Milan
  - ARM CPUs
    - *Fujitsu A64FX***
    - *Ampere ALTRA MAX**
    - *NVIDIA GRACE*
- RISC-V*

- 100+ users & 290+ accounts
- ~95 systems, mostly bare-metal
- Used by ATLAS, CMS, LHCb, QTI, ML research in IT department
- ~200 tickets handled p/a

E4 COMPUTER ENGINEERING

*Remote access via E4

SIMONS FOUNDATION

**Remote access via Simons foundation

Heterogeneous Architecture Testbed: Hardware

- RISC-V*
- CPU
  - Intel Xeon Gen 6, 5, 4 …
    - *Sierra Forest*
    - *Emerald & Sapphire Rapids*
  - IBM Power 8 & 9
  - AMD EPYC*: Genoa, Milan
  - ARM CPUs
    - *Fujitsu A64FX***
    - *Ampere ALTRA MAX**
    - *NVIDIA GRACE*
- Accelerators
  - NVIDIA Tesla GPUs
    - *L4, T4, A2**
    - *A100(X*), V100(S), P100, H100**
  - NVIDIA Bluefield 2 DPU
  - Intel GPUs
    - *Max 1100 (Ponte Vecchio)*
    - *Flex 140 & 170*
  - AMD
    - *MI100*, MI210x***
    - *ALVEO U200 (FPGA)*

- 100+ users & 290+ accounts
- ~95 systems, mostly bare-metal
- Used by ATLAS, CMS, LHCb, QTI, ML research in IT department
- ~200 tickets handled p/a

E4 COMPUTER ENGINEERING

*Remote access via E4

SIMONS FOUNDATION

**Remote access via Simons foundation

Heterogeneous Architecture Testbed: Hardware

- RISC-V*
- CPU
  - Intel Xeon Gen 6, 5, 4 …
    - Sierra Forest
    - Emerald & Sapphire Rapids
  - IBM Power 8 & 9
  - AMD EPYC*: Genoa, Milan
  - ARM CPUs
    - Fujitsu A64FX**
    - Ampere ALTRA MAX*
    - NVIDIA GRACE

- Networking Switches
  - InfiniBand
  - Intel Omni-Path

- Memory
  - Intel Optane

- Accelerators
  - NVIDIA Tesla GPUs
    - L4, T4, A2*
    - A100(X*), V100(S), P100, H100*
  - NVIDIA Bluefield 2 DPU
  - Intel GPUs
    - Max 1100 (Ponte Vecchio)
    - Flex 140 & 170
  - AMD
    - MI100*, MI210x**
    - ALVEO U200 (FPGA)
- AI specialised hardware
  - Cerebras*, Graphcore*

- 100+ users & 290+ accounts
- ~95 systems, mostly bare-metal
- Used by ATLAS, CMS, LHCb, QTI, ML research in IT department
- ~200 tickets handled p/a

*Remote access via E4

**Remote access via Simons foundation

Heterogeneous Architecture Testbed: Hardware

- **Euro HPC *Supercomputer* development access**
  - LUMI, Leonardo, Deucalion, Vega, etc …

- **Networking Switches**
  - InfiniBand
  - Intel Omni-Path

- **Memory**
  - Intel Optane

- **AI specialised hardware**
  - Cerebras*, Graphcore*

- RISC-V*

- **CPU**
  - Intel Xeon Gen 6, 5, 4 …
    - *Sierra Forest*
    - *Emerald & Sapphire Rapids*
  - IBM Power 8 & 9
  - AMD EPYC*: Genoa, Milan
  - ARM CPUs
    - *Fujitsu A64FX***
    - *Ampere ALTRA MAX**
    - *NVIDIA GRACE*

- **Accelerators**
  - NVIDIA Tesla GPUs
    - *L4, T4, A2**
    - *A100(X*), V100(S), P100, H100**
  - NVIDIA Bluefield 2 DPU
  - Intel GPUs
    - *Max 1100 (Ponte Vecchio)*
    - *Flex 140 & 170*
  - AMD
    - *MI100*, MI210x***
    - *ALVEO U200 (FPGA)*

- 100+ users & 290+ accounts
- ~95 systems, mostly bare-metal
- Used by ATLAS, CMS, LHCb, QTI, ML research in IT department
- ~200 tickets handled p/a

E4 COMPUTER ENGINEERING

*Remote access via E4

SIMONS FOUNDATION

**Remote access via Simons foundation

CERN openlab

AI AT CERN

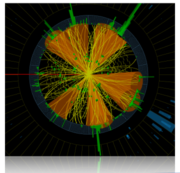

ML-*particle flow*

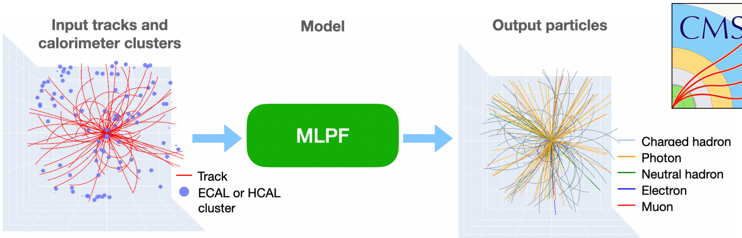

Accelerator systems

- Beam dynamics and control
- Enhanced diagnostics & predictions

Infrastructure

- Network saturation prediction
- Data quality monitoring

Real-time data selection & filtering

- O(nanos -micros) event reconstruction in FPGAs
- Fake reduction
- Anomaly detection w/ VAEs & CNNs
- Muon tracking
- On-detector data compression (AE)

Simulation

- Generative models for event generation and fast simulation e.g ATLAS FastCalo GAN

LHCb

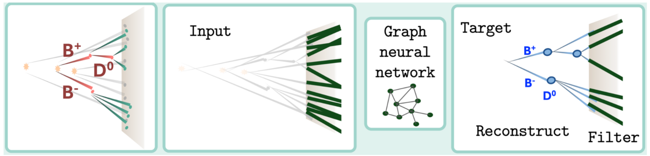

ATLAS EXPERIMENT

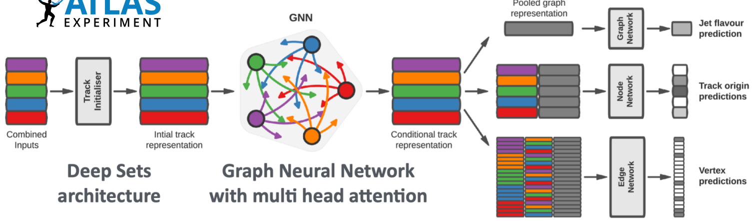

Analysis

- Clustering and pattern recognition
- Signal/background discriminations e.g *H-> γγ*
- ML-*particle flow*
- Particle classification & Jet tagging
- Energy calibration

CERN openlab

AI TRAINING AND INFERENCE OPTIMISATION ON HPC

- CERN openlab is involved in projects that are **implementing and optimising HEP AI/ML algorithms and workflows for HPC**
- To ensure usability and duplicability of efforts, **best practices** for scaling AI models on HPC are developed and **documented**
- **Euro HPC access** for development, benchmarking, and large-scale AI

Enables Machine Learning and AI algorithms and processing techniques

Opens the possibility for real-time interactive simulations (Digital Twins)

Burst/elastic resource scheduling

A path to optimize energy usage

New resources for processing

Requires technology migration and redesigning of applications

Encourages stronger engagement with industry, other science communities and the HPC computing community

Requires strategic planning and communication with the existing distributed computing community

Requires collaborating with HPC sites to develop common solutions to overcome technical challenges, leveraging on externally funded initiatives (EuroHPC, EC funded projects, industry,..)

AI ON EDGE DEVICES

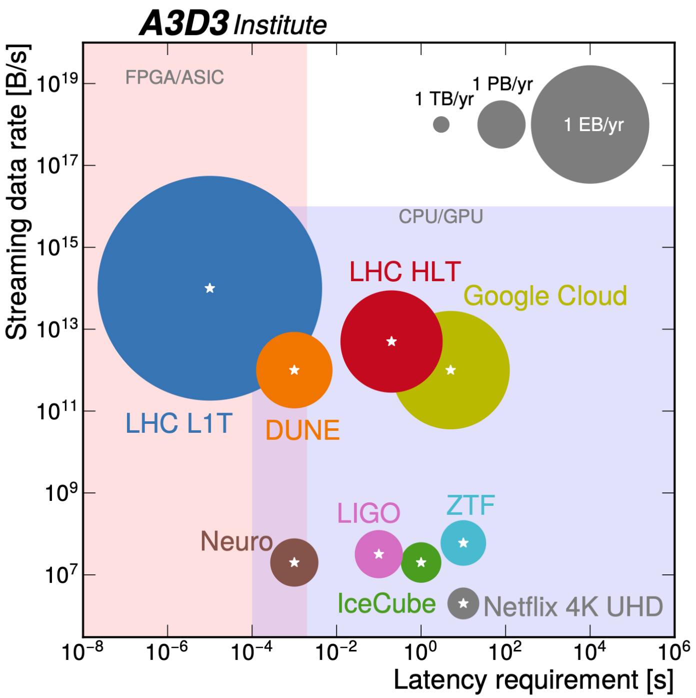

In collaboration with

LHC produces vast amounts of data every second billions of collisions per second during operation

- Without selection would generate ~ Pb/s raw data for CMS & ATLAS
- Impossible to readout/process/store all data
- Particles of interest rare among background
    - Need fast trigger to select interesting collisions for analysis with high efficiency, low fake rate

Fast ML *at the edge* needed for reducing and filtering data in real-time; ‘train offline’, ‘predict online’

Results from CMS and ATLAS demonstrate that ML on FPGAs can be used to improve selection efficiency and purity while keeping processing latency within fixed limits

https://cds.cern.ch/record/2852916?ln=en

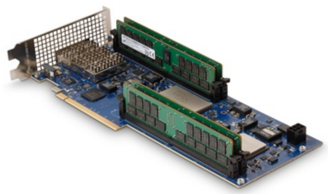

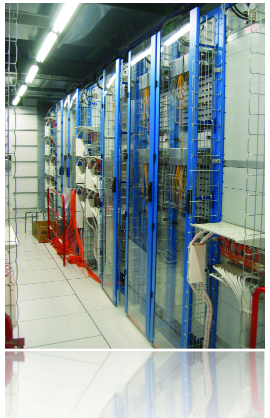

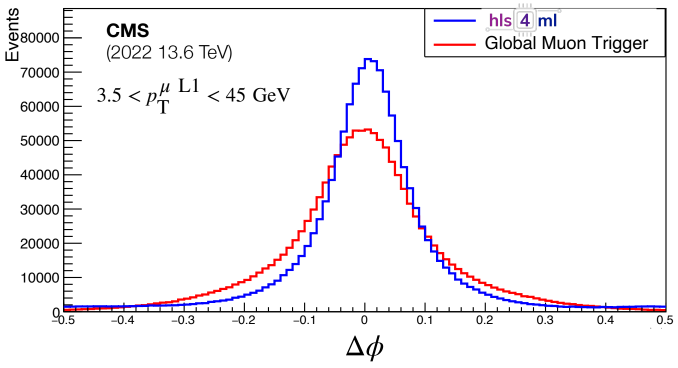

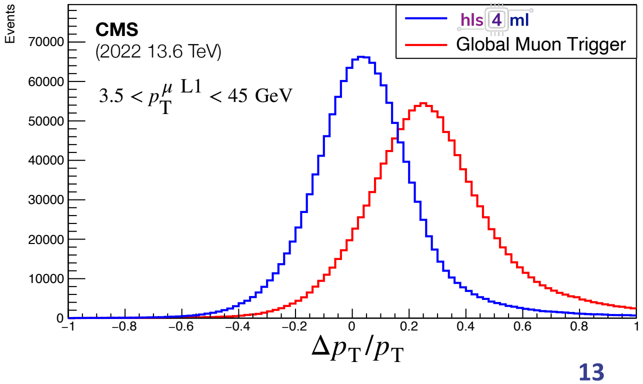

CERN openlab

REAL-TIME DATA PROCESSING ON CXL ARCHITECTURES

In collaboration with

micron

E4 COMPUTER ENGINEERING

CXL Compute Express Link™

- Near memory compute
- Low-latency
- High bandwidth
- Cache coherence

Emerging open standard for high-bandwidth heterogeneous, disaggregated computing

A memory lake architecture providing shared memory for heterogeneous computing units with a coherent view

memory lake **prototype for CMS data acquisition**

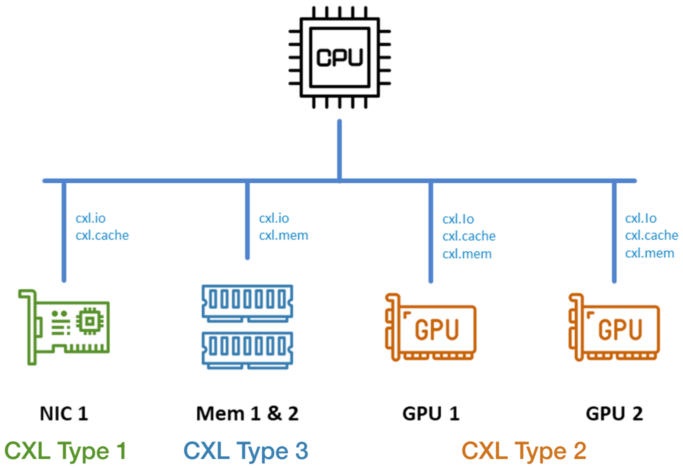

The three CXL device types and their protocols. Adapted from H3 (source)

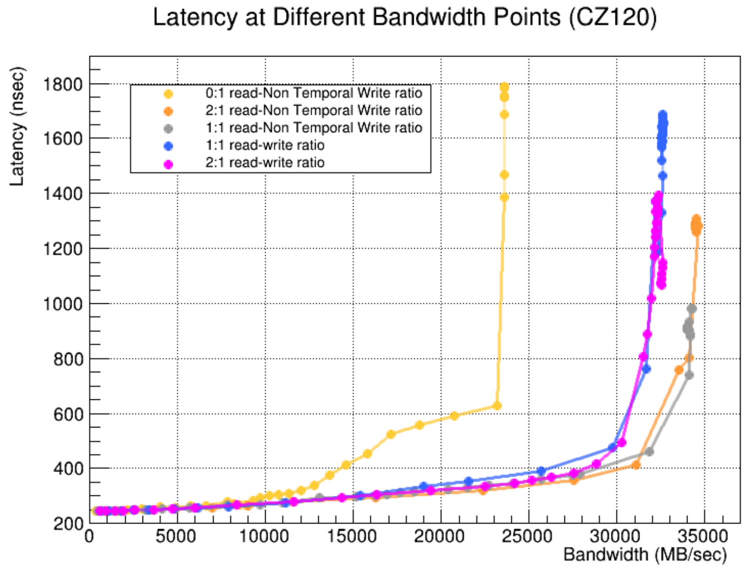

Micron CZ120 CXL memory modules

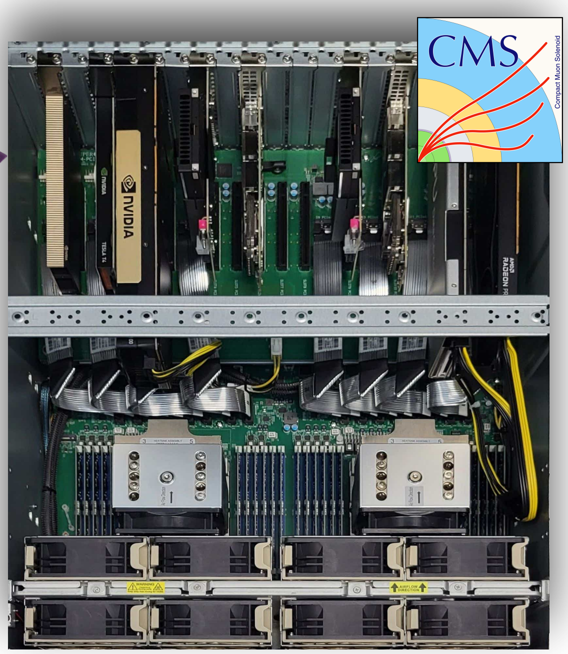

See poster on Thursday

CERN openlab

HYBRID AND CLOUD NATIVE APPLICATIONS

In collaboration with ORACLE

General move to container-based systems:

Cloud Native Computing Foundation:

- Open-source, vendor-neutral hub
- Hosting Kubernetes and its technologies landscape

Kubernetes:

- Open-source orchestration for automated deployment, scaling, management of containerised applications;
- CERN services already on platform (EDH, SSO, WLCG IAM, Gitlab, Rucio, SWAN, etc …)
- Advantages: interoperability (hybrid & multi-cloud); repeatability, resilience, deployment speed, auto-recovery
- Leveraging for sustainability, cost modelling, and disaster recovery [link]
- Evolving towards cloud-native AI

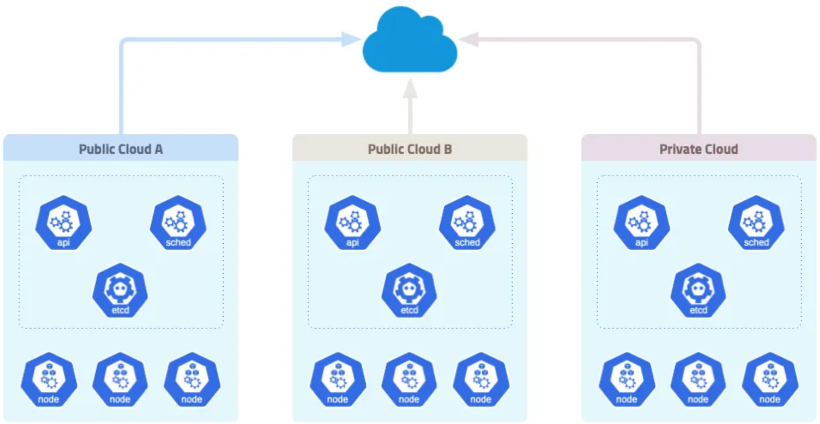

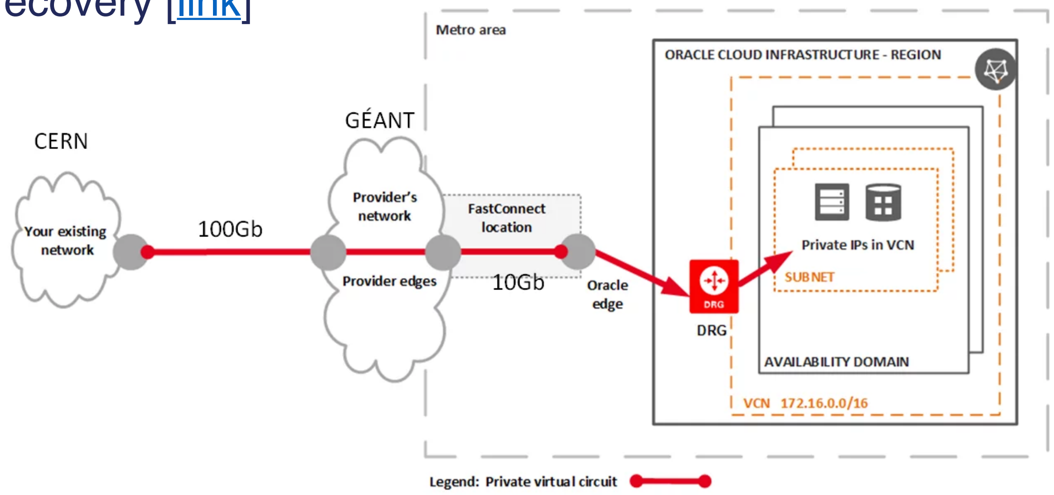

[Direct network interconnect between GÉANT backbone and the Oracle Cloud Infrastructure]

CERN openlab

NEW MATERIALS & TECHNOLOGIES FOR STORAGE SOLUTIONS

In collaboration with

cerabyte

PURESTORAGE®

- Evaluating emerging storage solutions
- Pioneering sustainable infrastructures for data storage and archiving

Magnetic Tape → **Ceramic Storage** / **Glass/Quartz Storage** / **DNA Storage**

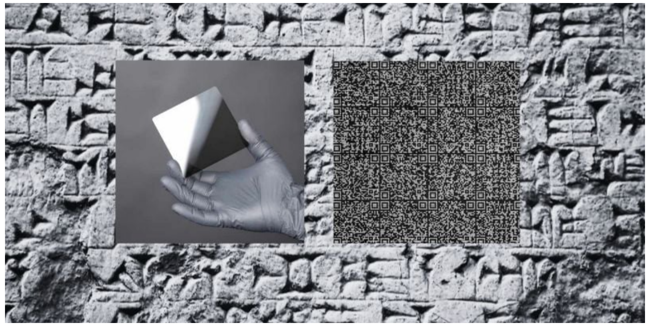

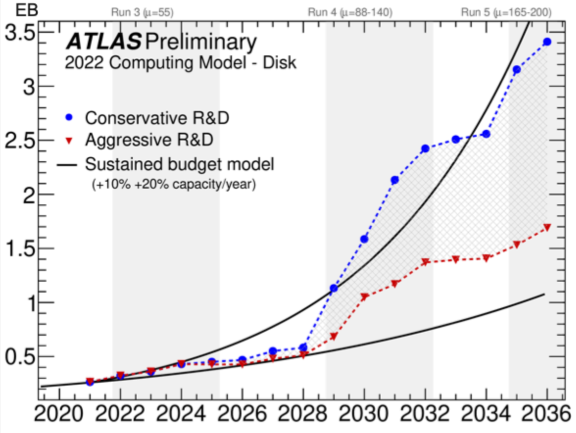

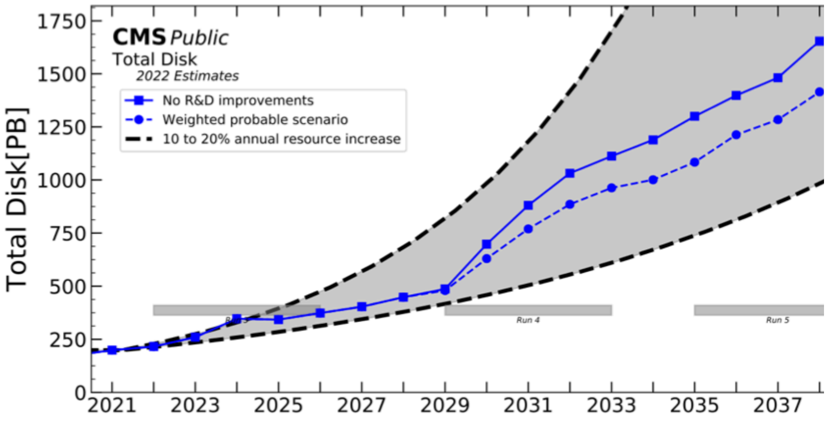

ATLAS and CMS online storage predictions

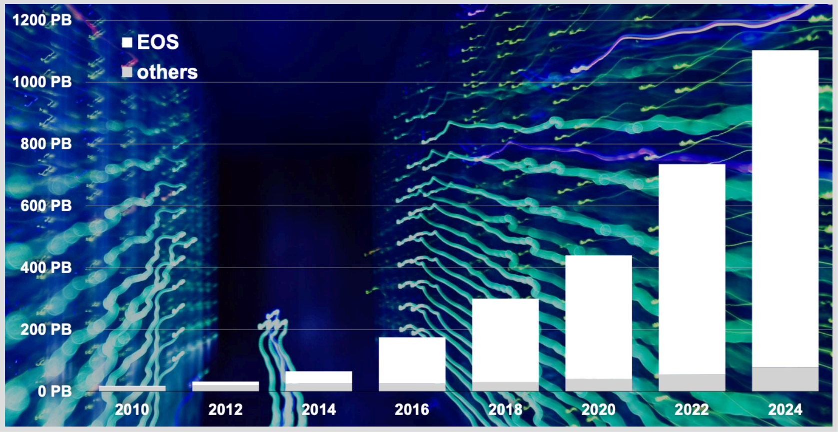

CERN IT - Operated Disk Storage Capacity

CERN openlab

COMMUNICATION, EDUCATION & OUTREACH

As a part of the education and training programme, CERN openlab **runs various initiatives that support participation of young scientists and other research organisations**

Summary Student Programme

Provides undergraduate and master's level students with an opportunity to work on one of the R&D projects for nine weeks under experts' supervision

This year there was a record of more than 6600 applicants!

Technical Workshop

Annual workshop to review the R&D projects carried out during the last year and discuss future plans. The event features technical talks, a poster session and a technology track dedicated to our industrial partners

Lectures & Training

Open access to CERN openlab **lectures that cover a wide range of computing topics**, from AI to exascale computing and quantum technologies. Regular **specialised technical training** to members of the scientific community

CERN openlab relies on Communication, Education & Outreach actions

SUMMARY

- CERN openlab has proven to be an invaluable mechanism by which to involve industry and access cutting-edge architectures in service of developing, testing, and integrating these new technologies.
- Meeting the future demands of the particle physics community requires embracing new hardware architectures and a heterogeneous hardware infrastructure.
- Openlab projects both diverse and impactful, across many domains.

Current members

SIEMENS · micron · E4 COMPUTER ENGINEERING · ORACLE · cerabyte · UNIVERSITÀ DEGLI STUDI DI TRIESTE · SIMONS FOUNDATION · INFN · intel

Phase VIII Brochure

Maria Girone
Head of Openlab

Thomas James
CTO for AI and Edge devices

Antonio Nappi
CTO for Platforms and Workflows

Luca Mascetti
CTO for Storage

Luca Atzori
CTO for Computing

Killian Verder
CTO Office Administration

Mariana Velho
Chief Communications Officer

Kristina Gunne
Chief Admin and Finance Officer

https://openlab.cern/

CERN openlab

DIGITAL TWINS OF ACCELERATORS AND DETECTORS

- Real-time virtual representation of physical object, system, or process
  - Mirroring its behaviour, performance, and environment
- Enables users to monitor, analyse, and optimise physical counterpart by interacting with the digital model
- Rely on AI/ML models to simulate the response
  - Cloud or on-prem. resources are sufficient for inference;
  - Large-scale training on extensive scientific datasets, hyper-parameter optimisation requires HPC.
- InterTwin 'unified digital twin engine' for AI-based scientific DTs

Uses @ HL-LHC: Detector prototyping and optimisation
  - Fast simulation of detector response to varying beam/environmental conditions

- Exploring NVIDIA OMNIVERSE for visualisation

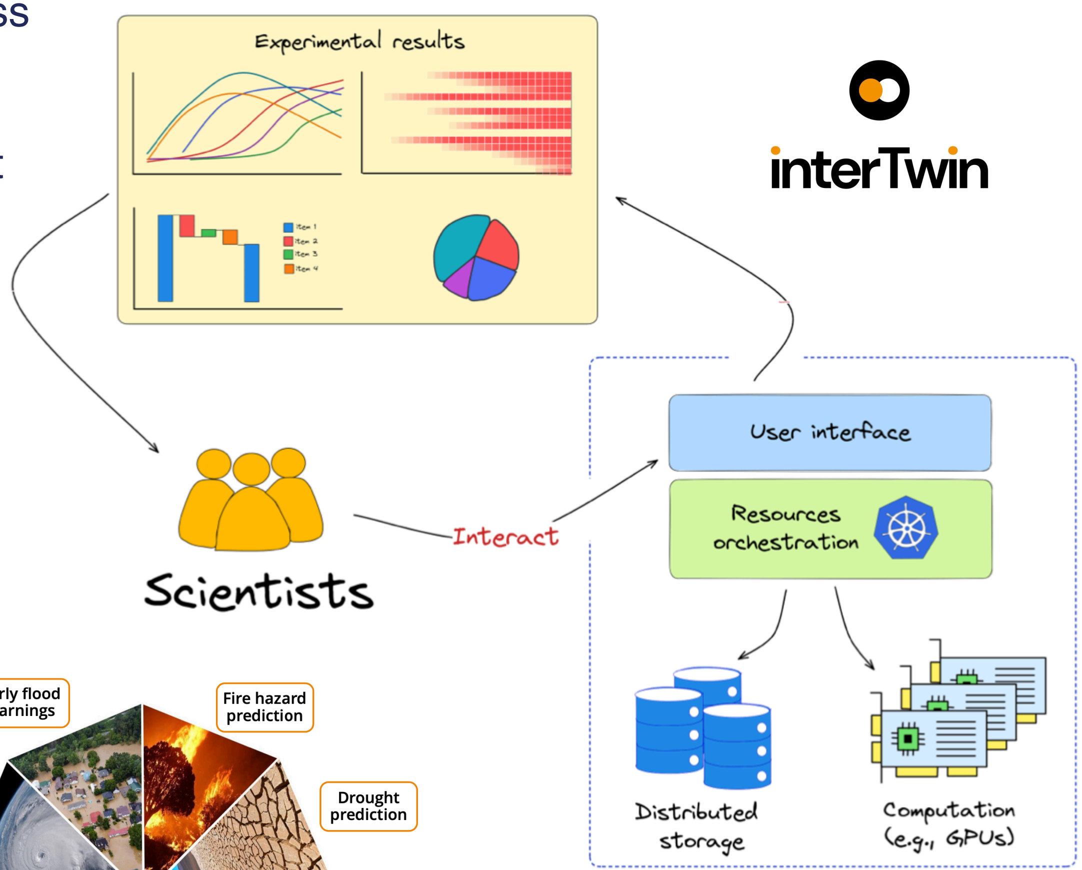

interTwin