



Annual Report

2025



Table of Contents

Foreword 1-2

A word on the future 3-4

CERN openlab 5-6

About CERN openlab 7-8

Vision & Strategy 9-10

Governance 11-12

Projects reporting for 2025 13-14

Anomaly Detection for Ultra Low Latency Event Selection at the LHC 15-18

Investigation of Anomaly Detection Algorithms for Filtering Events with Microseconds Latency in the ATLAS Hardware-Level Trigger 19-20

Expanding CERN’s Expertise in NVIDIA Technologies for Scientific Computing 21-22

Heterogeneous Architectures Testbed 23-24

Industrial Edge-Cloud and AI-based Agents 25-26

Next Generation Archiver for WinCC OA 27-28

Next-Generation Exascale Flash Storage 29-30

Evaluation of Cerabyte – Archival Data Storage Technology using Ceramic Nanolayers 31-32

Applied Multi-Disciplinary AI on High-Performance Computing 33-34

Oracle Kubernetes Operator 35-36

Cost Optimization and Sustainability for Public Cloud Provider (Carbon Aware FinOps) 37-38

Digital Twin: Data Science Engine 39-40

BioDynaMo: Biology Dynamics Modeller 41-42

Geometric Quantum Machine Learning with Neutral Atoms 43-44

SPECTRUM 45-46

ODISSEE 47-48

interTwin 49-50

Training & Education 51-52

Communication & Outreach 53-54

Foreword

by the Head of CERN openlab

Maria Girone, head of CERN openlab, photographed in the CERN Meyrin Data Centre.



2025 has been a remarkable and energising year for CERN openlab. Our mission, to accelerate computing for science through collaboration between CERN and industry, academia, and research organisations, has continued to grow in both ambition and scope. In a rapidly evolving technological landscape, CERN openlab remains a space where emerging ideas can be tested against the demanding and real-world requirements of frontier scientific research.

This year, we launched new projects, each addressing strategic technological challenges central to CERN's scientific programme. These projects span areas such as advanced computing architectures, AI-driven workflows, data management at scale, and next-generation networking, all essential to supporting the evolving needs of the LHC experiments and the wider scientific community. Together, they reflect our commitment to exploring technologies that will underpin scientific discovery in the coming decades.

Beyond individual projects, 2025 has also reinforced the importance of CERN openlab as a platform for long-term, forward-looking collaboration. By working closely with our partners, we are able to anticipate future computing requirements, explore innovative solutions early, and reduce the risks associated with deploying new technologies at scale. This collaborative approach allows CERN and its partners to learn together, adapt quickly, and translate research outcomes into practical, deployable solutions.

We have also been delighted to welcome new partnerships that bring fresh expertise, new capabilities, and renewed energy to our innovation ecosystem. They strengthen our links with both academia and industry, enriching our projects with diverse perspectives and complementary skills. At the same time, we have opened many promising new dialogues with potential partners, expanding the horizon for future joint research and technology development.

The continued growth of our community demonstrates both the relevance of CERN openlab's model, and the shared enthusiasm of our partners to push the boundaries of what is technologically possible. Through this growing network, CERN openlab acts as a bridge between cutting-edge research and technological innovation, ensuring that advances made in support of particle physics can also deliver broader benefits to society.

Together, we are building solutions that not only empower CERN's scientific mission but also contribute to technological progress benefiting society at large; from more efficient data processing and sustainable computing infrastructures to advances that can be applied in fields such as industry, healthcare, and environmental research.

As we look ahead, the coming years promise exciting developments. With the HL-LHC era approaching and digital technologies advancing faster than ever, the opportunities for impactful innovation are immense. CERN openlab will continue to play a key role in preparing CERN's computing ecosystem for this future, while remaining open to new ideas, new partners, and new ways of working.

In 2026, we will celebrate 25 years of CERN openlab, a milestone that reflects the strength, resilience, and vision of this unique public-private partnership. It is an opportunity not only to look back at what we have achieved together, but also to look forward with confidence and ambition.

I look forward to continuing this journey with our partners and collaborators, embracing new challenges and shaping the future of scientific computing through collaboration, innovation, and curiosity.

A word on the future

from the CERN Director for Research & Computing

Gautier Hamel de Monchenault, CERN Director for Research & Computing, photographed in CERN's building 40.



As CERN prepares for the High-Luminosity LHC era, the role of advanced computing has never been more central to our scientific mission. The HL-LHC will require innovation in computing technologies, infrastructures, and methodologies due to the unprecedented data volumes, complexity, and performance requirements.

For over twenty-five years, CERN openlab has provided a unique framework for collaboration between CERN, industry, and academia, enabling the co-development and early evaluation of emerging technologies in a real scientific environment. Through its projects and partnerships, CERN openlab contributes directly to strengthening CERN's computing ecosystem, while also fostering knowledge exchange and innovation that extend beyond particle physics.

Looking ahead, CERN openlab's focus on emerging technologies and sustainable infrastructures is well aligned with CERN's long-term strategy. By exploring new computing architectures, advanced data management solutions, artificial intelligence, and energy-efficient infrastructures, CERN openlab is helping to lay the foundations for future scientific discoveries. It is equally important to note its role in developing skills, building expertise and nurturing the next generation of computing specialists.

As CERN moves towards the HL-LHC and beyond, collaboration will remain essential. CERN openlab is a prime example of how close cooperation with industry and research partners can accelerate innovation and deliver solutions that benefit both science and society. I am confident that CERN openlab will continue to be a key enabler of CERN's scientific programme.

Beam screens for the High-Luminosity LHC.



CERN openlab

Fast, precise and groundbreaking science relies on powerful computing technologies, like the advanced algorithms and next-generation architectures that sift through the LHC's torrents of data in real time. At CERN openlab, new digital solutions are developed and tested under the demanding conditions of frontier research. From AI systems pushed to their limits to data platforms operating at unprecedented scale. The insights gained here fuel the know-how needed to drive discovery forward and to shape the future of scientific computing.

About CERN openlab

CERN openlab is a unique public-private partnership that brings together CERN and leading technology companies, research institutes, and universities to accelerate innovation in computing for science. Through collaborative R&D projects, it provides an environment where cutting-edge technologies can be tested, adapted, and advanced to meet the extreme demands of particle physics. By uniting expertise from industry and academia with CERN’s real-world challenges, CERN openlab helps drive breakthroughs that benefit not only the global scientific community, but society at large.

The CERN openlab team photographed in the CERN Meyrin Data Centre.



Since its establishment in 2001, CERN openlab has stood at the forefront of technological innovation as a unique public-private partnership. Created to bridge CERN with leading companies in the ICT sectors, CERN openlab provides a collaborative framework to address the exceptional computing challenges arising from CERN’s scientific mission. By working closely with industry and academic partners, it enables cutting-edge technologies to be tested, refined, and optimised in a highly demanding, research-driven environment.


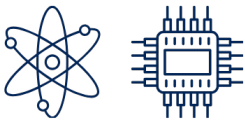




Over the years, CERN openlab has evolved into a dynamic innovation ecosystem, playing a role in meeting the computing needs of CERN’s experiments while supporting major scientific breakthroughs. Its collaborative projects have explored a wide range of advanced technologies, including high-performance computing, artificial intelligence, and scientific digital twins, helping to ensure that CERN remains at the leading edge of computational science.

CERN openlab operates through structured three-year phase cycles, designed to systematically evaluate technological developments, anticipate future requirements, and define strategic research priorities. This approach ensures a forward-looking and relevant R&D programme, while providing a stable framework for long-term collaboration and innovation. Its strategy is built around two R&D directions: Emerging Technologies and Sustainable Infrastructures.

Membership in CERN openlab offers access to a unique ecosystem characterised by unparalleled computing challenges, world-class scientific research, and a concentration of leading experts. This environment provides an ideal testbed for emerging technologies, allowing partners to validate and demonstrate their solutions through realistic and highly demanding use cases. In many cases, this close collaboration leads directly to measurable improvements in products, services, and technical capabilities.

CERN openlab’s long-term success lies in its ability to bring together global expertise and frontier technologies to address complex challenges at the intersection of science and computing. Looking ahead, it will continue to serve as a catalyst for collaboration, innovation, and excellence, ensuring that CERN remains a global leader in both scientific discovery and technological advancement.

CERN openlab at a glance

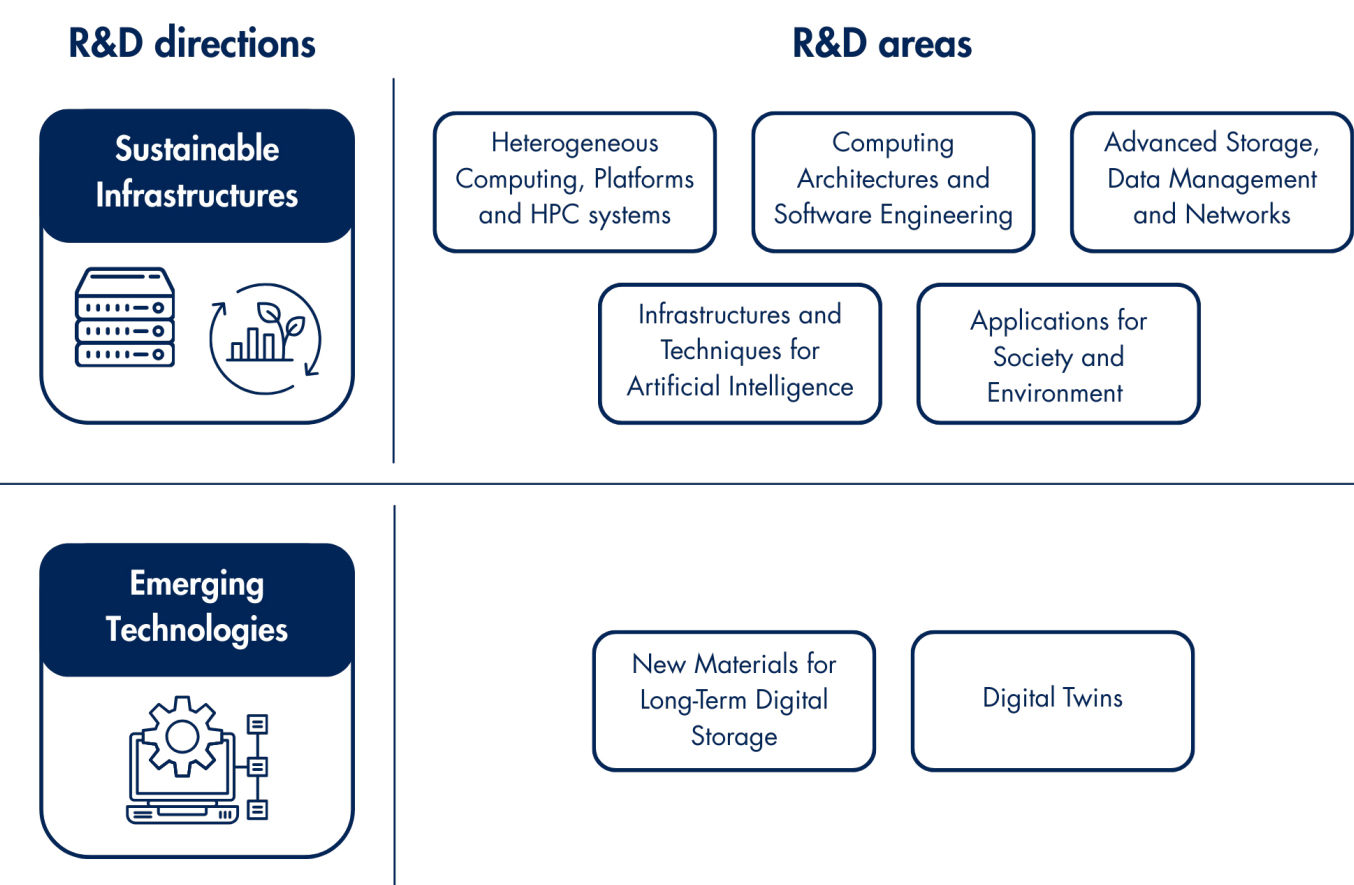
Stakeholders  CERN • Industry • Research Institutions	Mission  Advancing computing for science	Objectives  Test, refine, and optimise cutting-edge technologies
Challenge  CERN’s demanding computing requirements	Our projects  Structured collaborative R&D	R&D Directions  Emerging Technologies • Sustainable Infrastructures

Vision & Strategy

CERN openlab’s vision is to advance scientific computing through collaboration with industry and academia, ensuring CERN remains ready for the data-intensive challenges ahead. Its strategy is built around two R&D directions: Emerging Technologies, which explores innovations such as AI and new computing architectures; and Sustainable Infrastructures, which focuses on energy-efficient, scalable platforms for long-term scientific needs. Together, they provide a clear roadmap for driving technological progress that supports CERN’s mission while promoting sustainability.



R&D Directions



Throughout its history, CERN has been at the forefront of big data scientific research, with CERN openlab playing a pivotal role in tackling the associated computing challenges. By fostering collaborations with industry and research organisations, CERN openlab empowers the HEP community in its research endeavours. In response to the evolving landscape of scientific research, including the advent of exascale computing, CERN openlab spearheads efforts to enhance and scale up IT infrastructure to tackle the upcoming data challenges.

This collaborative approach drives innovation and advancement for all parties involved. Moreover, it enables partners to leverage solutions from HEP to address challenges in other fields, ensuring maximal relevance and impact. Projects within the CERN openlab framework are dedicated to accelerating computing for science, particularly under the following R&D directions and areas.

Governance

The governance of CERN openlab is built around a central coordination team, supported by relevant IT functions at CERN, to ensure efficient oversight and guidance. Membership is structured in two levels with Partners having access to the CERN openlab Collaboration Board and playing a leading role in long-term co-development programmes, whereas Associates engage in more focused, tactical projects. This governance model balances flexibility and structure: it supports a diverse community of industry, academic, and research institution members while ensuring that projects align with CERN’s computing needs and long-term R&D strategy.



Team



Maria Girone
Head of CERN openlab



Antonio Nappi
CTO for Platforms & Workflows



Luca Atzori
CTO for Computing



Luca Mascetti
CTO for Storage



Eric Wulff
CTO for AI on HPC



Thomas Owen James
CTO for AI and Edge Devices



Valentina Clavel
Finance Manager




Mariana Velho
Communication Manager

Members



Projects Reporting for 2025

CERN openlab projects address a wide range of computing and technological challenges in support of CERN's scientific programme. In 2025, these projects advanced collaborative R&D activities with industry and academic partners, delivering results, insights, and lessons learned across diverse areas. The following pages provide an overview of the activities carried out during the year, together with an outline of next steps that will build on the work completed and help shape future developments.



Anomaly Detection for Ultra Low Latency Event Selection at the LHC

ATLAS Experiment

This project aims to implement anomaly detection algorithms on FPGA devices in order to look for physics that may have been missed with standard trigger techniques. From the ATLAS side, the goal is to implement various anomaly detection algorithms in one of the FPGAs of the existing L1Topo trigger and validate them using LHC data from 2026. The knowledge and experience acquired during this process will be of use for the Upgrade stage of the project, which will consist of developing anomaly detection algorithms for the High Luminosity LHC. The CMS and ATLAS experiments are working together to share experience and pitfalls, making use of this solution already during the Run-3 at the LHC.

Project Coordinator

Stefano Veneziano

Technical Team

Paula Martinez Suarez
Stefano Veneziano
Ralf Gugel (Mainz University)

Overview

Extensive use of AD algorithms is foreseen in High Luminosity LHC. This work will allow to test some first prototypes using real-time LHC data, before the Long Shutdown 3 starting next year. They are being deployed in the ATLAS hardware trigger, built on FPGAs and designed to work at a very low latency (the latency envelope requirement is about 25ns. This effort will provide the operational experience needed to fully profit from this technology in the future.

Highlights in 2025


From the ATLAS side, during 2025 all the software and firmware tools used to train, translate, synthesise and implement the model in custom hardware have been tested and well understood.

These include TensorFlow (doi.org/10.5281/zenodo.4724125) and Keras (keras.io) for floating-point training, QKeras for quantisation-aware training (github.com/google/qkeras), hls4ml for user-friendly Python to HLS translation (doi.org/10.5281/zenodo.1201549), AMD Vitis™ HLS to generate the RTL design (amd.com/en/products/software/adaptive-socs-and-fpgas/vitis/vitis-hls) and AMD Vivado™ for synthesis and implementation in the AMD (Xilinx) FPGA (amd.com/en/products/software/adaptive-socs-and-fpgas/vivado).


A first model was deployed this year and is currently in operation in the ATLAS Experiment data flow. A second model has been developed in parallel to test possible enhancements and is planned to be deployed for the 2026 data-taking period.

Next Steps

The next steps include completion of performance and validation studies of the model deployed in 2026 and start contributing to the design of a new wave of anomaly detection algorithms within the frame of the High Luminosity upgrade, in the LOGlobal subsystem in particular (<https://atlas-l0global-docs.web.cern.ch/gbl-overview/index.html>).




The ATLAS Level-1 Topological module, containing two ultra-low latency FPGA-based processing units.



The ATLAS Level-1 Topological trigger subsystem installed in the ATLAS data acquisition room. One of the boards is hosting our Anomaly detection algorithm.

Publications & Presentations

S. Veneziano, Anomaly Detection for Ultra Low Latency Event Selection in ATLAS (4 March 2025). Presented at CERN Openlab Technical Workshop, Geneva, 2025.



Anomaly Detection for Ultra Low Latency Event Selection at the LHC

CMS Experiment

This project targets the deployment of real-time anomaly detection algorithms on FPGA devices to identify physics phenomena that may evade standard, model-driven trigger selections in the CMS experiment. The central objective is the implementation of a machine-learning architecture based on a transformer model, adapted for ultra-low-latency inference in hardware. This work is being carried out in collaboration with the ATLAS experiment, allowing both experiments to share their experiences.

Project Coordinator

Thomas Owen James

Technical Team

Elias Leutgeb
Thomas Owen James
Fabian Helmberger

Collaboration Liaisons for the ATLAS & CMS project with Oracle

John Lathouwers
Ludovico Caldara
Sengul Chardonnerau
Audrey Poulin
Cris Pedregal
Garret Swart
Arno Schneuwly
Dan Tow

Overview

The HL-LHC era will significantly increase data rates and event complexity, making traditional, model-dependent trigger strategies increasingly restrictive. Anomaly detection offers a powerful, model-agnostic alternative for identifying rare or unexpected physics signatures in real time. This project investigates whether modern heterogeneous hardware platforms, specifically AMD Adaptive Compute Acceleration Platforms (ACAPs), can meet our strict latency, throughput, and determinism requirements. These devices integrate programmable logic, processors, and Adaptive Intelligence (AI) Engines, offering a promising path toward scalable, real-time ML-based triggering and scouting solutions for CMS and future experiments.

Highlights in 2025

During 2025, the proposed transformer-based architecture was designed, simulated, and functionally verified using the AMD Vitis unified software platform. The fundamental building blocks of the transformer, most notably the matrix multiplication units, were explicitly implemented and aggressively optimized through tiling strategies, dimensionality constraints, and resource-aware sizing. Computationally intensive softmax operations were approximated using low-order polynomial functions to significantly reduce resource usage while preserving acceptable numerical performance.


Additional components, including concatenation layers and input/output buffering, were developed and optimized for execution on AMD AI Engines. Particular attention was given to the synchronization and scheduling of multiple AI Engine tiles, addressing dataflow and coordination challenges in order to maximize throughput and sustain real-time operation. This work also builds upon and adapts existing firmware developments from the CMS community, especially those originating from the Next Generation Trigger project. These designs were further optimized to satisfy the specific latency, resource, and bandwidth constraints imposed by the anomaly detection use case, providing a realistic path toward integration into CMS trigger and scouting environments.

Next Steps

The next phase of the project focuses on deploying the optimized Transformer architecture onto AMD hardware. This will require further adaptation of the model and its implementation to fully exploit the target architecture and meet strict real-time constraints. In parallel, the choice of model architecture will be revisited, potentially requiring simplifications or structural changes, followed by retraining using updated and hardware-aware input features. Performance, latency, and stability studies on real hardware will determine whether deployment is feasible within the CMS L1 trigger or 40 MHz L1 scouting system.

Publications & Presentations


J. Pearkes, CMS Experience with Anomaly Detection in the L1 Trigger. Presented at CERN openlab technical workshop, Geneva, 2025.



The current Global Trigger crate, which makes the final selection decision at L1 based on traditional and anomaly-detection algorithms.



The technical. From left to right Elias Leutgeb, Fabian Helmberger, Thomas Owen James.



Investigation of Anomaly Detection Algorithms for Filtering Events with Microseconds Latency in the ATLAS Hardware-Level Trigger

This project aims to evaluate the robustness of a candidate data-driven autoencoder-based anomaly detection algorithm for use at the ATLAS hardware trigger, which performs real-time event selection to save only the events deemed most interesting. This work tests the sensitivity of the algorithm across a variety of physics processes and explores new methods for increasing sensitivity to anomalous events occurring at low energy or low multiplicity, which are more easily obscured by background.

Project Coordinators

Paula Martinez Suarez
Stefano Veneziano

Technical Team

Paula Martinez Suarez
Ioannis Xiotidis

Collaboration Liaisons

Nikos Konstantinidis
Noah Clarke-Hall

Overview

The ATLAS detector has a data throughput of $O(\text{PB/s})$, far too high for persistent storage of the full read-out, necessitating real-time event selection and substantial data reduction via the Trigger and Data Acquisition system. The trigger reduces the event rate by a factor of 400, admitting only the “most interesting” events for detailed analysis. What the trigger defines as “interesting” directly modulates the physics accessible to the ATLAS collaboration, making it vital to ensure that events signalling new physics are not systematically discarded.

Highlights in 2025

Data analysis was carried out to quantify biases in the autoencoder algorithm with respect to L1 trigger variables, such as particle multiplicities and energies.


Further analysis quantified the sensitivity of the autoencoder to 11 different interaction types, representing a first step towards understanding what the model classifies as “anomalous.”

A preliminary implementation of a post-hoc debiasing algorithm was developed and subjected to initial bias tests. This approach aims to connect what the autoencoder identifies as anomalous, given its training data, with what is considered physically interesting once known data biases are taken into account.


Next Steps

The preliminary implementation of the post-hoc debiasing algorithm requires more extensive validation, which will be carried out in early 2026. As the project is exploratory in nature, further studies will investigate alternative or complementary approaches to maximising the robustness of data-driven hardware trigger algorithms, with a view towards potential adaptation into future trigger development.


This project is part of a 6-month collaboration with UCL Centre for Data Intensive Science and Industry.



ATLAS Central Trigger Processor (CTP) in 2016.



ATLAS Central Trigger Processor (CTP) in 2016.



Expanding CERN’s Expertise in NVIDIA Technologies for Scientific Computing

Advance scientific research and IT capabilities through cutting-edge computing technologies. This partnership, developed within CERN IT’s Heterogeneous Architecture Testbed, focuses on three key areas: digital twins, advanced computing architectures, and high-precision numerical libraries.

Project Coordinator

Luca Atzori

Technical Team

Luca Atzori
Maria Girone
Matteo Bunino
Joaquim Santos
Albane Carcenac
Jessy Sobreiro
Stefan Roiser
Ioannis Xiotidis
Stefano Veneziano
Thorsten Wengler

Collaboration Liaisons

Martin Roosen
Tom Gibbs
Peter Messmer
Sebastian Kalcher
Adam Thompson

Overview

By leveraging NVIDIA’s expertise, CERN aims to enhance simulation workflows, optimize data processing, and improve computational accuracy, all of which are essential for high-energy physics research.

Highlights in 2025

The project has just started, and we are currently laying the foundation for its full development.

In 2025, CERN openlab co-organised two dedicated workshops in collaboration with NVIDIA, focusing on accelerator-based computing and AI technologies relevant to scientific research.


Next Steps

In 2026, directions include:


- Benchmarking GPUs/DPU’s for Monte Carlo/track reconstruction efficiency.
- Mixed-precision FP64/FP16 optimization on AI hardware for 10x gains in event generation and pile-up mitigation.
- NVIDIA Holoscan enabled low-latency GPU streaming pipelines at CERN for real-time detector data processing.
- Omniverse digital twins with CAD models and live streams.

Publications & Presentations

A. Koehler, NVIDIA Compute Update and Recent Developments in HPC Software. Presented at the CERN openlab Technical Workshop, Geneva, 2023



CERN openlab summer students at a workshop during the 2025 CERN openlab summer student programme.



Heterogeneous Architectures Testbed

This project aims to provide a diverse hardware portfolio for comprehensive technology testing. Focused on assessing the efficacy of various architectures, this initiative aims to provide valuable insights into the practical utility of emerging technologies. By subjecting a spectrum of hardware configurations to real-world applications, the project seeks to establish benchmarks that guide the adoption of the most effective and efficient technologies.

Project Coordinator

Luca Atzori

Technical Team

Luca Atzori
Albane Carcenac
Maria Girone
Joaquim Santos
Jessy Sobreiro
David Southwick
Eric Wulff

Collaboration Liaisons

Cosimo Gianfreda
Marco Cicala
Daniele Gregori
Ian Fisk
Claudio Bellini
Bruno Riva
Walter Riviera
Arndt Mitwer
Frank Kuypers
Jean-Laurent Philippe

Overview

The project is pivotal for CERN and broader scientific communities. By creating a diverse portfolio of applications tailored for various hardware architectures, the project aims to enhance technology evaluation. This work is crucial for optimizing computational efficiency, fostering innovation, and ensuring that CERN and the wider scientific community stay at the forefront of technological advancements, ultimately advancing our capabilities in high-performance computing and scientific research.

Highlights in 2025

The project has achieved significant milestones in evaluating cutting-edge technologies across diverse CPU platforms, including the latest x86 (Intel® Xeon® 6th gen, AMD EPYC™ 9004 series) and Arm (Ampere® One®, NVIDIA Grace™) CPUs as well as GPUs such as Intel GPU Max and NVIDIA H100, both on premises and remotely through cloud resources.

The project extensively ran benchmarks with a focus on energy consumption measurements, providing a comprehensive assessment of efficiency. Notable applications and benchmarks such as HEPsScore23 were thoroughly tested.

Additionally, the testbed allowed extensive software testing for the CMS, ATLAS, and LHCb experiments. The support continuously given to the user community emerged as a crucial aspect, possibly the most important, fostering collaboration, offering valuable feedback, and ensuring seamless integration of new hardware.

2025 also represents the first step for this project towards the integration of external resources such as remote data centers, HPC centers or public commercial cloud resources.

This overall progress played a pivotal role in fostering collaboration with industry partners, offering valuable feedback on technology performance and shaping future developments.

The project's outcomes are integral not only for advancing CERN experiments but also for guiding industry stakeholders in optimizing their technologies for real-world applications. The collaborative and open approach ensures seamless integration of new hardware, enhancing computational capabilities and fostering groundbreaking advancements in scientific research.

Next Steps

The project's next steps involve ongoing evaluation of emerging technologies, crucial for readiness in the High-Luminosity LHC era. Continuous assessment of advancements in x86, Arm and RISC-V CPUs, as well as upcoming GPU platforms, remains a priority.

On top on this, the investigations related to the potential integration of external HPC centers and commercial cloud resources in our testbed will continue in 2026. This forward-looking strategy ensures that the "Heterogeneous Architectures Testbed" stays at the forefront of technological innovation to meet the computational demands of the evolving scientific landscape.

Publications & Presentations

Sobreiro, J., Carcenac, A., Atzori, L., Santos, J. (2025). Heterogeneous Architectures Testbed [Poster]. CERN openlab Technical Workshop, Meyrin, Switzerland.

Atzori, L. (2025). The CERN openlab Heterogeneous Architectures Testbed (HAT) Project [Conference presentation]. CERN openlab Technical Workshop, Meyrin, Switzerland.

Southwick, D. (2025). Benchmarking and Energy Studies with HEPsScore [Conference presentation]. CERN openlab Technical Workshop, Meyrin, Switzerland.



Industrial Edge-Cloud and AI-based Agents

Profiling of Siemens’ virtual PLCs real-time capabilities. Design and implementation of a test bed for large scale distributed deployment exploration of vPLCs. Implementation of source code generation AI agents for PLC languages and integration with WinCC OA tooling. POC analysis and implementation of AI-enabled applications, integrating with Industrial Control Systems and Operation processes.

Project Coordinator

Fernando Varela Rodriguez

Technical Team

Filippo Berto
Fernando Varela Rodriguez
Abhit Patil

Collaboration Liaisons

Christian Kern
Silvio Becher
Stefan Langer

Overview

vPLC are a recent advancement that focuses on merging the fields of OT and IT by integrating the high reliability of PLCs with the convenience of IT infrastructures. Migrating PLCs applications to commercial off-the-shelf hardware allows to more easily provide maintenance to operation sites, removing complexities linked to the physical hardware devices, especially in distributed systems.

The thoughtful integration of AI agents in development and operational processes can substantially reduce the completion times of repetitive and non central tasks, leaving more time to the users to concentrate on their main activities.

Highlights in 2025

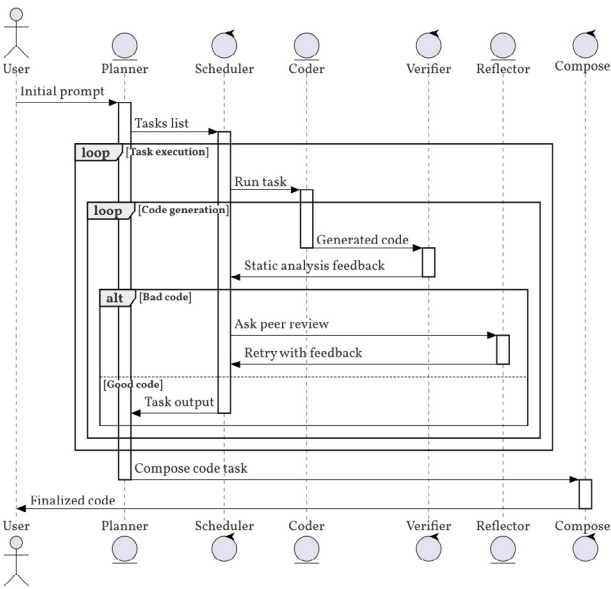
A dedicated edge-cloud testbed is being established at CERN to profile virtual PLCs across industrial edge devices, OpenShift, and local servers. This multi-tiered setup simulates real world conditions for comprehensive performance evaluation, focusing on behaviour, resource consumption, and bottlenecks. KPIs like cycle time, memory usage, I/O throughput, latency, and error rates are being tracked using profilers and monitoring agents.

A queryable knowledge base, built from WinCC OA and internal documentation, centralizes PLC programming and system architecture information. This resource facilitates efficient access to PLC function blocks, communication protocols, and troubleshooting guides. The knowledge base is being integrated with AI agents to explore automated PLC code generation using techniques like LLMs and code synthesis. This aims to accelerate development, reduce errors, and enhance code quality.

Furthermore, a use case analysis identified potential AI applications for DevOps and control room operations. Opportunities include predictive maintenance, anomaly detection, automated deployment pipelines, and real-time workflow optimization. This analysis prioritizes efficiency, reliability, and safety, guiding future AI development efforts and performance evaluation.

Next Steps

The immediate next steps involve executing comprehensive profiling experiments on virtual PLCs deployed across industrial edge devices, OpenShift, and local servers. A key focus will be the evaluation and benchmarking of Open Source Virtual Switch Solutions to explore their applicability for vPLC infrastructure. This assessment will determine whether OSS options can meet the performance, real-time, and reliability requirements, and if not, identify the specific gaps and improvements needed to enable their use in industrial deployment.



Code generation flow and actors.



Next Generation Archiver for WinCC OA

The project aims to modernize archiving of data from CERN’s WinCC OA-based SCADA systems by delivering a flexible and future-proof solution capable of handling growing data volumes and evolving use cases. It improves the performance and accessibility of historical data, reduces the dependence on a single database technology, and empowers engineers and scientists with richer tools for analysis and operational insight.

Project Coordinator

Rafal Kulaga

Technical Team

Martin Zemko
Nikita Nekhotyachshiy
Rafal Kulaga
Ewald Sperrer
Pedro Agostinho

Collaboration Liaisons

Christopher Stoegerer
Ewald Sperrer
Pedro Agostinho

Overview

WinCC OA is the de facto standard for developing SCADA systems at CERN, supporting over 850 mission-critical systems across experiments, infrastructure, and accelerators. Growing data rates and increasing needs for data retrieval, both within WinCC OA and from external systems, make performance and functional improvements essential. The NextGen Archiver addresses these challenges as a modular and future-proof solution with pluggable backends supporting multiple databases. At CERN, it will be used with Oracle for backward compatibility and TimescaleDB for improved query performance and new use cases.

Highlights in 2025

Significant progress on the TimescaleDB schema, backend, and test suite was achieved in 2025. The solution’s performance has been validated at the scale of large production deployments, with hundreds of simultaneous ingestion and querying processes. In all scenarios, it met or exceeded the requirements, even on moderately performant hardware.

To enable simple and efficient retrieval of historical data without requiring SQL knowledge, a dedicated API has been developed. Together with Grafana extensions – initially developed as part of an openlab Summer Student project – that allow users to browse and select signals using familiar table and tree views, this functionality addresses a long-standing need for creating rich and dynamic dashboards. These dashboards will simplify the daily work of hundreds of system experts, operators, and scientists at CERN.

To unlock the full benefits of the new archiving solution, data from existing Oracle schemas must be migrated to TimescaleDB. This represents a significant challenge, not only due to the scale – nearly 200 TB distributed across more than 350 schemas – but also because of the complex validation checks required to ensure full preservation of data integrity.

Next Steps

In 2026, the team will focus on completing the TimescaleDB support and preparing it for CERN-wide deployment in the first half of 2027. Production-grade versions of Grafana extensions and data retrieval APIs will also be finalized based on pilot user feedback, alongside the transfer of 200 TB of existing data from Oracle to TimescaleDB.

Publications & Presentations

A. Kveton, R. Kulaga, M. Zemko - “Optimizing time series data storage for CERN industrial control systems using TimescaleDB and PostgreSQL” Presented at CERN PGDay 2025, Geneva, 2025, <https://indico.cern.ch/event/1471762/contributions/6280212/>



Next-Generation Exascale Flash Storage

This project aims to evaluate next-generation, high-density flash-based storage technologies through a strategic CERN openlab - Pure Storage collaboration. By combining CERN's exascale operational expertise with Pure Storage's high-efficiency DirectFlash platform, the initiative will assess performance, scalability, energy efficiency, cost, and reliability. The overall goal is to determine whether such technologies can sustainably and cost-effectively support future scientific data volumes at exabyte scale.

Project Coordinator

Luca Mascetti

Technical Team

Andreas J. Peters
Robert-Paul Pasca
Ruhi Choudhury

Collaboration Liaisons

Ethan Miller
Franck Moreau
Hari Kannan
Anthony Glidic
Kevin Kremer

Overview

CERN's experiments, particularly in the High-Luminosity LHC era, will generate data volumes that exceed the capabilities of today's storage architectures. To sustain scientific productivity, CERN must evaluate technologies that can deliver far greater density, performance, and energy efficiency at scale. Industry advances in NAND-flash provide a unique opportunity to rethink how exabyte-level systems are built and operated. By collaborating with Pure Storage, CERN can assess disruptive methods that may substantially reduce cost, power consumption, and operational complexity. The outcome has the potential to benefit not only CERN but also the broader scientific community confronting similar data-intensive challenges.

Highlights in 2025

The project has made significant progress in evaluating Pure Storage's FlashBlade S500 appliance as a high-performance backend for large-scale physics workflows. Extensive benchmarking was completed using MPI-based distributed tests, IO500 workloads, and EOS integration scenarios, providing a multidimensional view of throughput, latency, scalability, and metadata behavior on both NFSv3 and NFSv4.1 protocols. Testing confirmed that the FlashBlade delivers strong read and write performance, peaking near 75 GiB/s and 24 GiB/s respectively under IO500 on NFSv3, while also highlighting protocol-specific bottlenecks, caching effects, and sensitivity to concurrency levels.



The CERN openlab and Pure Storage teams in the CERN Meyrin Data Centre.

A custom MPI benchmarking tool was developed to measure real-time EOS transfers via XRootD, enabling controlled experiments across uploads, downloads, and mixed patterns. Integration work also addressed a key compatibility gap by implementing an EOS plugin that emulates extended attributes, allowing FlashBlade to be integrated within EOS. Tiering simulations further demonstrated the performance gap between flash and HDD layers, validating the strategic value of flash acceleration in hierarchical storage models.

Overall, these initial results indicate strong potential for flash-based storage within CERN workflows, while also identifying areas requiring tuning, such as network utilization, protocol behavior, and future support for parallel file access.


Next Steps

Next steps focus on targeted optimisation and deeper integration testing. Priorities include refining mount and protocol configurations, fully characterising NFSv4 behaviour, validating asynchronous XRootD performance at scale, and re-running benchmarks after upcoming FlashBlade firmware updates. Additional work will assess multi-tier EOS deployments, explore parallel file-access capabilities, and establish decision thresholds for production readiness.

Additionally, we plan to evaluate the performance of the PureStorage EXA design in the CERN environment and perform a comparative evaluation against the XRootD protocol.



Pure Storage technology installed in the CERN Meyrin data centre.



Evaluation of Cerabyte – Archival Data Storage Technology Using Ceramic Nanolayers

This project will allow CERN to gain understanding of the possibilities and potential problems of the new archive storage medium using ceramic nanolayers. It represents an opportunity to provide feedback and to guide the development of this technology in a way that is beneficial to CERN.

Project Coordinator

Vladimir Bahyl

Technical Team

Vladimir Bahyl
Mostafa Solitan

Collaboration Liaisons

Ed Childers
Sebastian Kirsch
Martin Kunze

Overview

Currently tape is CERN's archival storage medium of choice, due to its characteristics of reliability, durability and price. While the roadmap for tape is encouraging, tape as a technology is exposed to certain market risks, and CERN should always have an eye on the development of possible alternatives and likely evolution of archival storage.

For many years, IT-SD group has actively engaged with companies doing R&D in innovative storage technologies, including DNA, archive glass and ceramic. At this point, ceramic seems to be the most likely contender to achieve sufficient I/O rates to be useful, within a reasonable time frame (~10 years). The success of this technology would represent a disruptive change to the landscape.

Highlights in 2025

In 2025, CERN received first samples of ceramic glass data carriers from Cerabyte. These contained data encoded in QR codes. The purpose of the exercise was to understand the physical and noise characteristics of this new medium.

With the help of the Mechanical and Materials Engineering Group from the CERN Engineering Department we gained access to a digital microscope KEYENCE VHX-7000 and with 500x magnification we were able to read and decode the data from the QR codes.

During the summer, CERN hired a student to work on improving the speed of the QR detection to industrial scale. He managed to optimize the WeChat QR code decoding algorithm for GPU and reached around 18000 QR code detections per second.

In August 2025, CERN technical team visited Cerabyte's lab in Vienna, Austria.

Next Steps

As the project came to the end of the original 24-month agreement during 2025, both sides agreed on extending it for another 24 months. Once the legal aspects are clarified, the collaboration will continue by focusing on the following areas:

- Development of a reliability model for Cerabyte's read/write process by optimizing QR code selection, placement, and erasure coding as a function of the system's raw bit error rate characteristics.
- Development of an open, self-describing format to ensure the long-term reliability of data written with this technology.
- Demonstration of write and read of CERN data.

Publications & Presentations

M. Solitan: Optimizing QR Code Decoding for Cerabyte's Ceramic Data Storage (19 August). Presented at CERN IT-SD group meeting, Geneva, 2025





Image showing QR codes arranged in a matrix as stored on a Cerabyte data carrier.



Applied Multi-Disciplinary AI on High-Performance Computing

The project aims to enable scalable and future-proof machine-learning workflows for large-scale scientific data analysis, with a focus on high-energy physics event reconstruction. It addresses both computing and algorithmic challenges, optimising data access and storage for efficient training and inference, while developing machine-learned reconstruction methods that are accurate, scalable, and adaptable to evolving detectors and computing architectures.

Project Coordinator

Eric Wulff

Technical Team

Eric Wulff
Joosep Pata
Radomir Wasowski
Enrico Bocchi
Abhishek Lekshmanan

Collaboration Liaisons

Ian Fisk
Andras Pataki
Mariel Pettee

Overview

This project is composed by two use-cases: 'Ceph scaling strategies for machine learning workloads' and 'HPC-accelerated AI optimisation'. Machine learning applications continue to grow in complexity, driving increasing model sizes and unprecedented demands on data storage and access. While software-defined storage solutions, such as the POSIX-compliant Ceph File System, are well suited to large-scale, distributed ML workflows, storage access speed and coordination overheads remain key limitations. This project focuses on optimising Ceph resource usage to improve performance for large-scale training and inference. In parallel, it aims to modernise high-energy physics event reconstruction by replacing hand-crafted heuristic algorithms with scalable machine-learned approaches. Building on the Machine-Learned Particle Flow (MLPF) framework, the project also explores supervised, self-supervised, and foundation-model techniques to deliver accurate, efficient, and adaptable reconstruction suited to future detectors and computing architectures.

Highlights in 2025

At the start of the project, reliable benchmarking of the storage system was established. Dedicated hardware was provisioned to measure raw performance, while supporting software layers were benchmarked to quantify overheads introduced at each level of the storage stack and to control error sources. Under controlled synthetic workloads, both Ceph's internal metrics and client-side I/O measurements were used to study the impact of configurable system parameters. In parallel, detailed performance-profiling capabilities were enabled to identify and analyse suboptimal code behaviour.

Methods to emulate storage-access patterns typical of machine-learning workloads, including those used at CERN, were systematically reviewed. Particular emphasis was placed on tools that minimise computational requirements, reducing interference with benchmarking results and dependence on hardware with limited availability. Suitable tooling was identified, addressing a key early project objective.

In 2025, the foundations for an end-to-end self-supervised learning workflow were also established, building on experience from Machine-Learned Particle Flow (MLPF). A shared repository and data-processing pipeline were developed, including dataset pre-processing tools and a full PyTorch-based training pipeline. Initial studies explored multiple self-supervised approaches, alongside supervised baselines for comparison.

Next Steps

This project will continue next year, characterizing the performance of more complex storage cluster configurations with a focus on horizontal scaling. Various ML workloads will be tested utilizing the benchmarking methodology developed earlier, their performance with conventional Ceph scaling strategies analysed to identify promising approaches, and profiling conducted to guide the development of optimization.

The next phase of the project will also focus on consolidating the effort on self-supervised learning and assessing its potential across a range of reconstruction-relevant tasks. Work will continue to explore and compare different clustering, SSL, and foundation-model approaches for learning meaningful representations directly from low-level detector data, while refining the data formats, workflows, and evaluation strategies.

Publications & Presentations

Wasowski, R. M., Bocchi, E. (2025, March 4). Applied Multi-Disciplinary AI on High-Performance Computing [Conference presentation]. 2025 CERN openlab Technical Workshop, Meyrin, Switzerland
<https://indi.to/f4nM8>

Wulff, E. (2025, March 4). Hyper-Parameter Optimisation on HPC systems [Conference presentation]. 2025 CERN openlab Technical Workshop, Meyrin, Switzerland.
<https://indico.cern.ch/event/1440389/contributions/6364347/>



Oracle Kubernetes Operator

This project aims to use the Oracle Kubernetes Operator and other Cloud Native Computing Foundations tools to modernise the Oracle Database and Oracle REST Data Services (ORDS) services at CERN. By implementing the Oracle Operator, the CERN team aim to modernize database lifecycle management through native Kubernetes automation. This transition will streamline provisioning and configuration into manageable, declarative resources. Furthermore, it bridges the gap for CERN ORDS infrastructure to fully adopt GitOps workflows. Additionally, the Oracle Operator will complement tools like Crossplane, supporting the implementation of Infrastructure as Code for infrastructure management.

Project Coordinator

Antonio Nappi

Technical Team

Antonio Nappi
Artur Wiecek
Miroslav Potocky
Thomas Saury

Collaboration Liaisons

John Lathouwers
Marco Stefanetti
Matteo Malvezzi
Ruggero Citton
Cristobal Pedregal

Overview

CERN operates a large number of Oracle databases supporting critical scientific, administrative, and data services. As CERN continues its transition towards cloud native platforms, it is essential to modernise how these databases and related services, such as ORDS — used to provide REST APIs for Oracle databases — are deployed and managed. This project evaluates the Oracle Kubernetes Operator and CNCF tools to automate database and ORDS lifecycle management, improve security and consistency, and enable GitOps based workflows. The outcomes aim to reduce operational complexity at CERN, and would be applicable to other research institutions and Oracle users pursuing similar modernisation efforts.

Highlights in 2025

In 2025, significant progress was made on integrating Oracle REST Data Services (ORDS) with Kubernetes in close collaboration with Oracle teams. CERN actively tested the Oracle Database Operator and the OrdsSrvs controller, identifying functional gaps, configuration limitations, and GitOps-related challenges. Multiple issues were reported and discussed with Oracle engineers, leading to concrete fixes and design improvements.

Key advancements include enhanced secret management for ORDS, notably the introduction of wallet-based authentication using Kubernetes Secrets, support for mounting Oracle wallets and TNS configuration files, and the ability to manage multiple database pools securely. CERN also contributed feedback on CRD (Custom Resource Definition) design, duration settings, password handling, and init scripts, which influenced upcoming releases of the operator.

In parallel, CERN evaluated and successfully prototyped ORDS deployment using the Central Configuration Server model, identifying its strengths and limitations compared to operator-based deployments. This work resulted in new examples, documentation updates, and feature additions on Oracle’s development branch, including Central Configuration support in OrdsSrvs.

Overall, the collaboration led to tangible upstream improvements, clearer deployment patterns, and a more secure and flexible ORDS architecture aligned with CERN’s Kubernetes and GitOps requirements.

Next Steps

In 2026, the project will focus on validating the new OrdsSrvs features introduced in the upcoming Oracle Database Operator 2.1 release, particularly Central Configuration support and wallet-based authentication. In parallel, CERN plans to evaluate the Oracle Operator for managing selected Oracle 26ai databases deployed on-premise. Further work will assess observability (OTel), production readiness, and GitOps alignment to define a sustainable ORDS deployment model.

Technical contributions to open-source projects

Technical contributions through issue reporting and participation in technical discussions, supporting improvements to operator functionality and reliability (GitHub issues [#186](#) and [#188](#)).



Cost Optimization and Sustainability for Public Cloud Provider (Carbon Aware FinOps)

This project aims to identify and implement metrics to understand the carbon footprint of the current CERN services hosted on Kubernetes. These metrics will inform the deployment of workloads in function of the Scope 3 emissions across on-premises and public clouds as a path to achieve low/zero-carbon workload deployment. Ultimately, these innovative results will have direct impact in the IT Department cloud operating model, where carbon awareness will be a costing factor being considered in the FinOps activity managing public cloud costs and optimizing cloud spending within CERN.

Project Coordinator

Ricardo Rocha

Technical Team

Laura Eve Sarah Llinares
Ricardo Rocha

Collaboration Liaisons

John Lathouwers
Dan Tow

Overview

New trends in computing show a tendency to rely on AI, ML and other computing models that rely heavily on new generations of accelerator technologies. This has pushed costs and power requirements up significantly in many cases for maximum performance with alternative solutions focusing instead on cost and power efficiency.

As more computing workloads at CERN also start relying heavily on GPUs and other accelerators as well as ML, it is important to better understand the workload patterns to properly match the ideal infrastructure. With the speed of evolution in this new generation of architectures, it is also important to ensure the required flexibility to incorporate external resources seamlessly into CERN workloads.

Highlights in 2025

During 2025 the project was started with an initial focus on research, experimentation and preparation of the first two key deliverables.

The first includes a report on the available tools for cost and carbon footprint metric collection, including an evaluation of projects such as Kepler and OpenCost within the CNCF ecosystem. The second report defines and compiles the relevant metrics for cost and carbon efficiency, to be validated against industry standards.

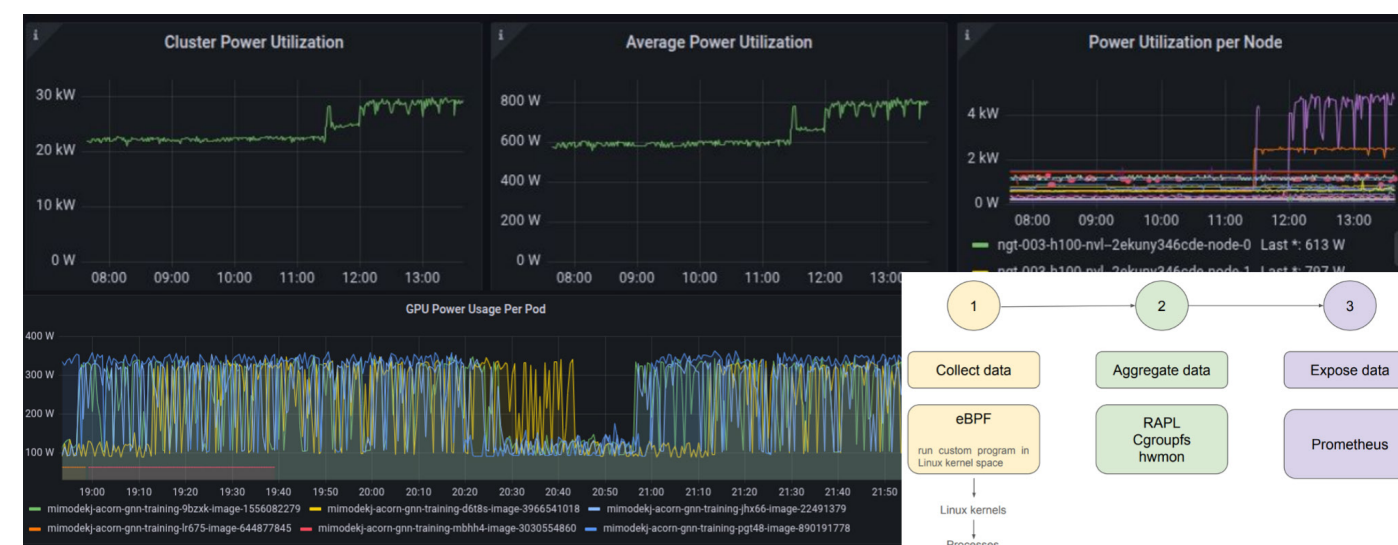
Practical tests were also performed by deploying Kepler and OpenCost on both Oracle and CERN cloud resources to begin collecting preliminary cost and energy usage data.

Next Steps


During 2026 the project will continue with the completion and validation of the report on cost and carbon efficiency metrics. The focus will then be on delivering a proof-of-concept deployment and implementation integrating the selected tools and generating metrics for reference use cases. The initial focus will be on reporting rather than automated scheduling, covering different types of workloads.

Publications & Presentations

Laura Llinares, Amine Lahouel, Overview of available power metrics. Presented to the Next Generation Triggers (NGT) Work Package 1.1 (05 September 2025): <https://indico.cern.ch/event/1580626/>



Cluster power observability for GPU workloads: end-to-end monitoring from kernel-level power collection (eBPF/RAPL), through cgroup-based aggregation, to Prometheus metrics visualized in Grafana at cluster, node, and pod level.



Digital Twin: Data Science Engine

This project aims to address critical gaps in data analysis and predictive modeling by employing advanced mathematical models and machine learning techniques for healthcare applications. A key objective is to develop a predictive engine based on these methodologies, ensuring higher accuracy and reliability in forecasting and trend analysis. The approach integrates cutting-edge optimization strategies to improve efficiency and effectiveness in complex systems. CERN will contribute its expertise in Digital Twin technology and advanced methodologies developed in the context of high-energy physics experiments.

Project Coordinator

Nicola Serra

Technical Team

Nicola Serra
Sara Sansaloni Pastor
Victorien Leconte

Collaboration Liaisons

Gang Mu

Highlights in 2025

In 2025, the project focused on strengthening collaboration and alignment between CERN and Johnson & Johnson. A series of dedicated meetings were held both at CERN and at Johnson & Johnson's offices in Basel, providing valuable opportunities to refine the shared vision, exchange expertise, and align expectations across scientific, technical, and organisational perspectives. These interactions helped consolidate the partnership, clarify priorities, and lay the groundwork for future technical activities by reinforcing mutual understanding and long-term collaboration.

Next Steps



Building on the strengthened collaboration established in 2025, the next phase of the project will focus on translating the shared vision into concrete technical activities. This includes identifying priority use cases, refining technical requirements, and exploring suitable data, models, and computing approaches to support the development of the Digital Twin Data Science Engine. These steps will prepare the ground for future implementation and experimentation.

Publications & Presentations

Mu, G. (2025, March 5). Digital Twin – Data Science Engine [Conference presentation]. 2025 CERN openlab Technical Workshop, Meyrin, Switzerland.



Maria Girone, head of CERN openlab, presenting at the meeting in the J&J offices in Basel.



Alveolar segment model in TOPAS-nBio from [Mechanistic model of radiotherapy-induced lung fibrosis using coupled 3D agent-based and Monte Carlo simulations](#).

BioDynaMo: Biology Dynamics Modeller

BioDynaMo is a platform that empowers scientists to effortlessly generate, execute, and visualize agent-based simulations. Utilizing cutting-edge computing technologies, the BioDynaMo platform facilitates simulations of unprecedented scale and complexity. This capability opens avenues for addressing intricate scientific research inquiries with greater ease.

Project Coordinator

Maria Girone

Technical Team

Maria Girone
Eric Wulff
Stavros Portokalidis

Collaboration Liaisons

Roman Bauer
Vasileios Vavourakis

Overview

In the life sciences community, computer simulation is gaining prominence for modelling intricate biological systems. While numerous specialized tools exist, creating a high-performance, versatile platform represents a significant advancement. CERN leverages its extensive expertise in large-scale computing, supported by funding through the CERN and Society Foundation, to address three highly relevant societal needs: the fight against cancer, inequalities and dengue. Additionally, BioDynaMo is part of the CERN IT department's efforts to build a Digital Twin Engine for Science, with its integration in projects such as interTwin, making it accessible to a wider range of researchers.

Highlights in 2025

In 2025, work focused on improving the usability and accessibility of the BioDynaMo platform. The installation process was streamlined to support a wider range of environments, and the documentation was expanded to assist new users and developers. These improvements enhance the platform's readiness for broader adoption within the research community. Integration with the interTwin Digital Twin Engine also advanced, enabling initial test simulations within the framework.

Progress on the tumour growth use case continued. A pancreatic tumour model that incorporates selected biological processes related to tumour progression and immune response was developed, guided by recent computational studies. The aim is to first replicate established computational tumour models to ensure consistency and validation, then develop a reproducible and extensible model that can support future research and biomedical insights.

BioDynaMo was also used in research by Cayla Harris et al. to study retinal development using agent-based modelling, illustrating how biological mechanisms and stochastic cellular decisions can be explored through simulation.

These developments contribute to strengthening the platform's reliability and usability while preparing it for a wider range of applications in digital biology and computational medicine.

Next Steps


The next phase will focus on completing and validating the pancreatic tumour model and documenting its implementation. Parallel work will continue to maintain and evolve the core software. A second biological use case will be initiated, simulating dengue and other infectious disease transmission dynamics, further broadening BioDynaMo's scientific applications.

Publications & Presentations

C. Harris, U. Abubacar, R. Bournes, T. Adel, A. Tamaddoni-Nezhad, R. Bauer, Agent-based modelling of retinal development. Bioinformatics and Biomedical Engineering (IWBBIO 2025), https://doi.org/10.1007/978-3-032-08452-1_6

S. Portokalidis BioDynaMo: Biology Dynamics Modeller (5 March). Presented at CERN Openlab Technical Workshop, Geneva, 2025. URL: <https://indico.cern.ch/event/1440389/>

S. Portokalidis BioDynaMo: Biology Dynamics Modeller (26 November). Presented at End-of-year Get together - CERN & Society Foundation, Geneva, 2025. URL: <https://indico.cern.ch/event/1615976/>



Geometric Quantum Machine Learning with Neutral Atoms

The collaboration will focus on neutral atom quantum computing techniques. The research collaboration will contribute to advance knowledge and tools that will be strategic to both Parties. The substantial innovation in the field of AI in the past years together with the rapid prototyping of quantum technologies, has enabled the definition of quantum machine learning. It is an active field of research which seeks to take advantage of the capabilities of both quantum computers and machine learning techniques, adapting the latter to the strengths of the former.

Project Coordinator

Michele Grossi

Technical Team

Jogi Suda Neto
Michele Grossi
Cenk Tüysüz
Andrea Gentile

Collaboration Liaisons

Jogi Suda Neto
Michele Grossi
Cenk Tüysüz
Andrea Gentile

Overview

Starting from a deep understanding of current classic and quantum implementation of theoretical and computational models for graph neural network, we will focus on scalability, symmetry properties and generalization. In particular, geometric deep learning (GDL) will be a key focus of this work package. The possibility of testing and implementing those model for HEP use cases represent an important test bed for future LHC computing requirements and quantum technologies.

Highlights in 2025

We were able to finalize the hiring of a PhD student and started the literature review part and setup of next goals.

Next Steps

The collaboration between Pasqal and the team at CERN QTI will transform the preliminary theoretical result into customized algorithms for High Energy Physics use case, taking advantage of the quantum platform developed by Pasqal.



Pasqal ORION quantum processor.



The SPECTRUM team at the all-hands meeting at CERN in January 2025.

SPECTRUM

The EU-funded SPECTRUM project unites top European science organizations and e-Infrastructure providers to develop a Strategic Research, Innovation, and Deployment Agenda (SRIDA) and a Technical Blueprint for a European compute and data continuum. This initiative aims to establish an Exabyte-scale research data federation and compute continuum, enhancing data-intensive scientific collaborations across Europe.

Project Coordinator

Sergio Andreatozzi (EGI)

CERN Technical Team

Maria Girone
Eric Wulff
David Southwick

Consortium Members

EGI Foundation
CERN
OCA
CNRS
INFN
FZJ
NEOVIA
SURF
CINECA
ASTRON



Spectrum is funded by the European Union Grant Agreement Number 101131550

Overview

Frontier research in High-Energy Physics and Radio Astronomy is entering the Exascale era, with new instruments requiring unprecedented data processing capabilities. To meet these demands, pan-European data and compute infrastructures must be developed, incorporating novel architectures, federation models, and IT frameworks. The integration of Exascale HPC and Quantum computing systems offers new opportunities to accelerate discoveries and complement existing research facilities. Key challenges include scalability, performance, energy efficiency, portability, interoperability, and cybersecurity, which must be addressed to ensure successful integration of these heterogeneous systems.

Highlights in 2025

In 2025, the SPECTRUM project has made excellent progress towards its main goals of creating a Technical Blueprint for a European Compute and Data Continuum as well as a Strategic Research, Innovation and Deployment Agenda (SRIDA). Work packages 3 and 5 finalized their work, started in 2024, consisting of a detailed study of HEP/RA use-cases, a detailed study on the RI infrastructure landscape, and conducting a comprehensive survey on current and future community needs. The work of WP3 and WP5 was reported in deliverables D3.1, D5.1, D5.2 and D5.3 and was submitted to and accepted by the European Commission (EC).

The development of the two most important outputs of SPECTRUM, the Technical Blueprint and the SRIDA began in January 2025 and has progressed rapidly. These documents are based on the insights gathered from surveys, use-cases, and SPECTRUM CoP discussions. Both documents were published as preliminary drafts in December of 2025 with the aim to gather feedback from the community. The publications will be followed by an organized consultation phase where the project gathers community feedback and shape the documents accordingly. The final version of the documents will be published in April of 2026.

Next Steps

In 2026, SPECTRUM will focus on finalizing and publishing the Technical Blueprint and SRIDA in April, incorporating feedback from the consultation phase. The remaining months until the project's conclusion in June will be dedicated to dissemination, community engagement, and ensuring the sustainability and adoption of SPECTRUM's outcomes across the European compute and data continuum ecosystem.

Publications & Presentations

- SPECTRUM D3.1: Community of Practice – Interim Report (2025). Zenodo. <https://doi.org/10.5281/zenodo.15100586>
- SPECTRUM D1.4: Policy Feedback Brief for Period 1 (2025). Zenodo. <https://doi.org/10.5281/zenodo.15100949>
- SPECTRUM D1.5: Dissemination, Communications and Exploitation – Interim Report (2025). Zenodo. <https://doi.org/10.5281/zenodo.15101144>
- SPECTRUM D5.1: Representative Use Cases – Analysis and Alignment (2025). Zenodo. <https://doi.org/10.5281/zenodo.15310223>
- SPECTRUM D5.2: Interoperable Access Policies – Analysis and Recommendations (2025). Zenodo. <https://doi.org/10.5281/zenodo.15647608>
- SPECTRUM D5.3: Landscape of Research Infrastructures – Technologies, Services, and Gaps (2025). Zenodo. <https://doi.org/10.5281/zenodo.15647755>
- Draft Technical Blueprint of a European Compute and Data Continuum (2025). Zenodo. <https://doi.org/10.5281/zenodo.17901518>
- Draft Strategic Research, Innovation and Deployment Agenda (SRIDA) (2025). Zenodo. <https://doi.org/10.5281/zenodo.17901518>
- Shaping the Future of Digital Infrastructures for Data-Intensive Science in High-Energy Physics and Radio Astronomy (2025). Zenodo. <https://doi.org/10.5281/zenodo.14772235>
- SPECTRUM Technical Blueprint and Strategic Agenda: Delivering Europe's Roadmap for Exabyte-Scale Scientific Infrastructure (2025). Presented at ACUD 2025. <https://pretalx.surf.nl/acud-2025/talk/QKLG7/>
- SPECTRUM SRIDA and Technical Blueprint Public Consultation Webinar (2025). <https://www.spectrumproject.eu/article/spectrum-public-consultation-webinar>



ODISSEE

ODISSEE aims to use AI to cope with the data deluge from SKAO and CERN’s HL-LHC by making their computing workflows more energy-efficient and sustainable. The project co-designs modular software and hardware building blocks, benchmarks emerging architectures, and improves workflow portability. This is an EU funded project that federates efforts from 3 pan-European ESFRI infrastructures (HL-LHC, SKAO and SLICESRI) in physical sciences, Big Data, and in the computing continuum supporting flagship instruments that will maintain and strengthen European leadership in high-energy physics and astronomy.

Project Coordinator

Damien Grataour
(Observatoire de Paris)

CERN Technical Team

Pierfrancesco Cifra
Francesco Sborzacchi
Niko Neufeld
Maria Girone
Matteo Bunino

Consortium Members

Observatoire de Paris	Simula
CNRS	SURF
CERN	GENCI
SKAO	EAR
ASTRON	SiPearl
INRIA	NextSilicon
BSC	NEOVIA
CSCS	EPFL

Overview

Next-generation instruments like the HL-LHC and SKAO will produce exabyte-scale data and significantly increase the power demand of computing and data-centre infrastructures. To cope with this challenge, the ODISSEE project aims to revolutionise the way we process, analyse and store data. ODISSEE is developing new technologies using AI to process and filter only the relevant data into the data stream on the fly. This approach will enable scientists to build more complex, yet reliable, physical models, whether at micro or astronomical scale.

At the same time, the project also focuses on redesigning hardware and software solutions that cover the entire data stream continuum, from generation to analysis. The primary objective is to create tools that are energy efficient, adaptable and flexible for the future. This will involve the development of a reconfigurable network of diverse processing elements driven by AI.

Highlights in 2025

In 2025 the focus was on preparing the digital-twin data-centre use case for later technical work. We started gathering requirements from the main scientific communities, including CERN’s LHCb experiment and CERN data centre. This work will prepare target scenarios for energy-aware operations and predictive maintenance, the relevant timescales, and acceptable trade-offs between energy savings, system utilisation and reliability.



We performed a survey of existing tools, frameworks and telemetry sources that could underpin the data-centre digital twin. This included studying solutions such as the Energy Aware Runtime (EAR) and related monitoring stacks, and mapping how they could expose power and performance metrics for both compute nodes and cooling equipment. Taken together, these activities provide a common vocabulary and technical baseline, and prepare the ground for implementing and validating a first digital-twin prototype on the LHCb data centre in the next project phase.

Next Steps

As a next step, we will implement and validate a first digital-twin prototype on the LHCb data centre. This will include integrating selected telemetry sources, instantiating power and reliability models for compute and cooling equipment, and running pilot optimisation and predictive-maintenance studies on a limited set of services. Results will guide further refinement and potential adoption at additional CERN and partner sites.



ODISSEE is funded by the European Union Grant Agreement Number 101188332



interTwin

interTwin seeks to co-design and implement a prototype of an open-source Digital Twin Engine (DTE), built upon open standards, facilitating seamless integration with application-specific Digital Twins (DTs). This innovative platform, rooted in a co-designed interoperability framework and the conceptual model of a DT for research, known as the DTE blueprint architecture, aims to simplify and accelerate the development of complex application-specific DTs.

Project Coordinator

Tiziana Ferrari (EGI)

CERN Technical Team

Matteo Bunino
Maria Girone
Eric Wulff
Kalliopi Tsolaki
Jarl Sondre Saether
Anna Lappe
Linus Maximilian Eickhoff
Enrique Garcia Garcia
Xavier Espinal
Sofia Vallecorsa

Consortium Members

CERFACS	EGI Foundation	LIP
CERN	EODC	MPG
CESNET	ETH Zurich	PSNC
CMCC	EURAC	TU Wien
CNRS	FZJ	UHEI
CSIC	GRNET	UNITN
Cyfronet	INFN	UPV
Deltare	JSI/IZUM	VU
DESY	KIT	WWU
ECMWF	KBFI	

Overview

interTwin develops and implements an open-source DTE that offers generic and customized software components for modeling and simulation, promoting interdisciplinary collaboration. The DTE blueprint architecture, guided by open standards, aims to create a common approach applicable across scientific disciplines. Use cases span high-energy physics, radio astronomy, climate research, and environmental monitoring. The project leverages expertise from European research infrastructures, fostering the validation of technology across facilities and enhancing accessibility. InterTwin aligns with initiatives like Destination Earth, EOSC, EuroGEO, and EU data spaces for continuous development and collaboration.

Highlights in 2025

In 2025, within interTwin we extended our itwinai-based integration from Juelich (JUWELS system) and Vega to additional EuroHPC systems: LUMI with AMD GPUs and Deucalion with Fujitsu A64FX CPUs, accessed via CERN openlab. On these platforms we benchmarked distributed AI workloads, studying strong and weak scaling through epoch times, per-iteration timing and detailed profiles of communication overhead between workers. We complemented this with measurements of GPU energy use, GPU and CPU utilisation, using them as signals to guide performance tuning and configuration choices. Hyper-parameter optimisation was exercised at scale to optimize the accuracy of AI-based digital twins. Finally, we demonstrated cloud-HPC integration using interLink, offloading distributed training jobs from a cloud environment to Vega HPC and performing scaling tests in the interTwin context to validate portability of our AI workflow stack.



interTwin is funded by the European Union Grant Agreement Number 101058386

Next Steps

The interTwin project has finished in 2025. Following the completion of the interTwin planned activities in 2025, the project has delivered valuable technical insights. The results will inform future R&D directions, projects, and provide a foundation for new initiatives and continued collaborations.

Publications & Presentations

M. Girone & M. Bunino, itwinai: Enabling Scalable AI Workflows on HPC for Digital Twins in Science (P29 poster). Presented at PASC25 – Platform for Advanced Scientific Computing Conference, Brugg, Switzerland, 2025.

M. Girone & M. Bunino, Enabling Lattice QCD Normalizing Flows in HPC Infrastructures (P110 poster). Presented at PASC25 – Platform for Advanced Scientific Computing Conference, Brugg, Switzerland, 2025.

M. Bunino, M. Girone, A. Lappe, J. Sondre Saether, Co-designing and Prototyping Interdisciplinary Digital Twins (15 min). Presented at 2025 CERN openlab Technical Workshop, CERN, Geneva, 5 March 2025.

M. Bunino & D. Ciangottini, Testing AI Containers for Digital Twins in Science: A Cloud-HPC Workflow. Presented at KubeCon + CloudNativeCon Europe 2025, 4 April 2025.

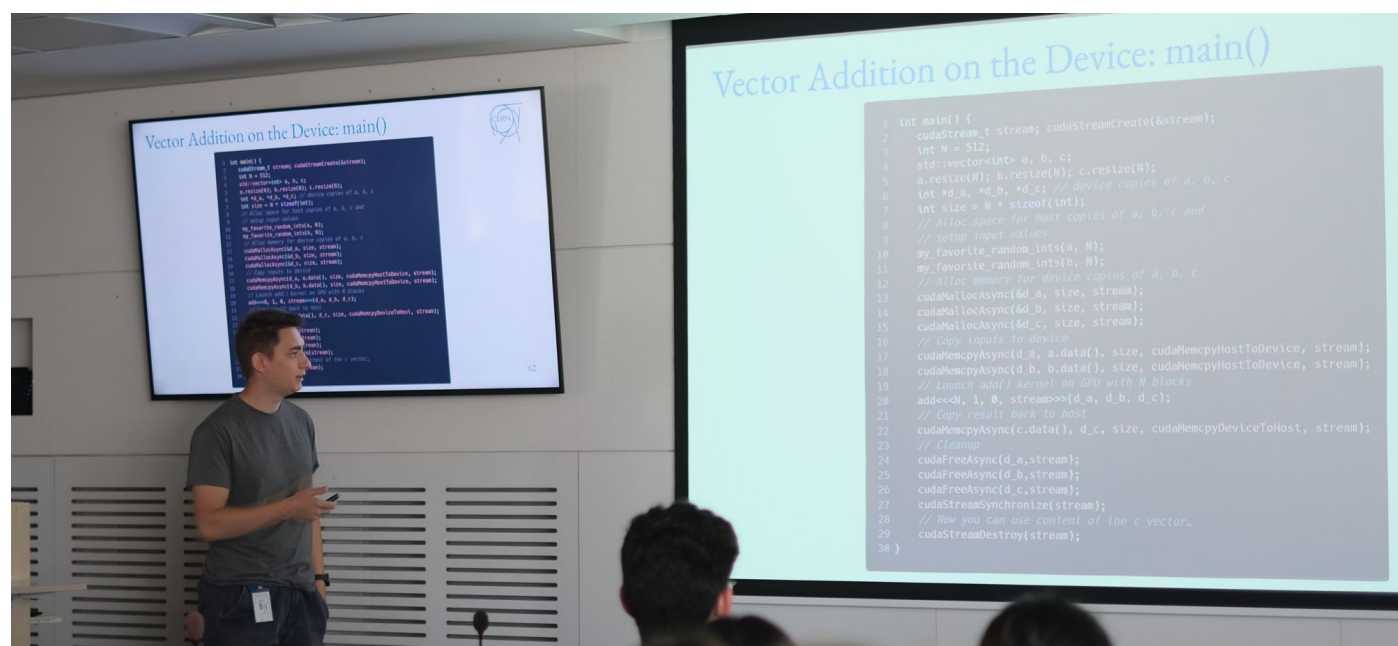
M. Girone & M. Bunino, Shaping the Future of Scientific Digital Twins: Tales from Physics, Climate Research, and Data Centre Operations (BoF session). Presented at ISC High Performance 2025, Hamburg, 11 June 2025.

Training & Education



CERN openlab's training and education activities aim to develop skills and expertise at the intersection of cutting-edge computing and scientific research. Through hands-on projects, workshops, lectures, and collaborative programmes, CERN openlab provides students, early-career researchers, and professionals with practical experience in advanced digital technologies. These activities support knowledge transfer, foster talent development, and help build a skilled community prepared to address the evolving computing challenges of modern science.

To prepare for the future of scientific computing, it is essential to ensure that the computing specialists of tomorrow have the skills needed to fully exploit emerging technologies. Through projects, lectures, and workshops, CERN openlab equips future computer scientists with essential knowledge that supports scientific progress and drives innovation. As part of its education and training programme, CERN openlab runs a range of initiatives that encourage the participation of young scientists and researchers from across the global research community.



Simone Balducci giving a training on the 'Fundamentals of Accelerated Computing with CUDA'.

Workshops & Hackathons

In 2025, CERN openlab organised and contributed to multiple workshops and hackathons that provided hands-on learning opportunities for both the wider CERN community and CERN openlab summer students. A series of technical workshops were organised through CERN openlab, covering topics such as advanced computing architectures, artificial intelligence, data management, and emerging technologies, and were open to researchers, engineers, and students alike. In addition, CERN openlab co-organised two dedicated workshops in collaboration with NVIDIA, focusing on accelerator-based computing and AI technologies relevant to scientific research. CERN openlab also promoted and supported participation in the [Open Quantum Institute](#) hackathon on quantum materials for sustainable development, in which CERN openlab summer students actively took part. Together, these activities fostered skills development, interdisciplinary collaboration, and practical engagement with cutting-edge technologies.

Collaborating with the Next Generation Triggers

In 2025, CERN openlab strengthened its training and education activities through close collaboration with the [Next Generation Triggers](#) (NGT) project within the CERN openlab summer student programme. A total of eleven students worked on NGT-related projects over a nine-week period, gaining hands-on experience with advanced trigger concepts, data processing, and computing challenges relevant to future experiments. This collaboration created valuable synergies between CERN openlab and the NGT community, while offering participating students enriched educational opportunities, exposure to real-world research environments, and direct engagement with cutting-edge developments in high-energy physics computing.

Partnering with ideas4HPC



In 2025, as in 2024, CERN openlab partnered with [ideas4HPC](#) to fund a female summer student, supporting the promotion of women in high-performance computing through targeted training and hands-on research at CERN. The collaboration also enables former funded students to present their work at the [PASC Conference](#), fostering professional development and visibility within the international HPC community.



Jarl Saether and Linus Maximilian Eickhoff teaching the 'Fundamentals of Deep Learning'.



2025 CERN openlab summer students photographed in front of building 513 (CERN Meyrin Data Centre).

CERN openlab Summer Student Programme

The CERN openlab summer 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. In addition, the public is granted access to CERN openlab lectures, which cover a broad spectrum of computing subjects, ranging from AI to exascale computing.

Lecture Programme

As part of the CERN openlab summer student programme, participants benefit from a dedicated lecture programme designed to complement their project work. The lectures cover a broad range of topics, including advanced computing technologies, data science, artificial intelligence, and high-performance computing. Delivered by experts from CERN, academia, and industry, the programme provides students with valuable theoretical context, exposure to current research challenges, and a broader understanding of the role of computing in modern scientific discovery.



Lightning Talks

As part of the CERN openlab summer student programme, lightning talks provide students with an opportunity to present their work in a concise and accessible format. These short presentations encourage knowledge sharing, communication skills, and peer-to-peer learning, while allowing students to gain experience in presenting technical topics to a broad scientific audience and to receive feedback from the CERN community.



Visits to CERN Facilities

Visits to CERN's facilities form an integral part of the CERN openlab summer student programme, offering students first-hand exposure to the experimental and computing infrastructure that underpins CERN's scientific mission. These visits help place students' project work in a broader context, deepening their understanding of how advanced computing supports large-scale experiments and reinforcing the connection between theoretical knowledge, practical implementation, and real-world scientific challenges.



Communication & Outreach

CERN openlab's communication and outreach activities are designed to support collaboration, increase visibility, and ensure the effective dissemination of results and expertise. A combination of dedicated events, conference participation, digital platforms, and targeted outreach initiatives is used to engage partners, researchers, and the wider scientific and technological communities. Through tools such as the CERN openlab Technical Workshop, its online presence, social media channels, and broader outreach activities, CERN openlab promotes knowledge sharing, community building, and sustained engagement around its research activities.

Technical Workshop

In 2025, the CERN openlab Technical Workshop continued to serve as a key forum for knowledge exchange and collaboration within the CERN openlab community. The workshop brought together partners, researchers, and experts from CERN, industry, and academia to share project updates, technical insights, and future perspectives across CERN openlab’s R&D activities. Through a combination of presentations, discussions, and interactive sessions, the workshop fostered dialogue, strengthened collaboration, and supported alignment on shared technological challenges and opportunities.

In 2025, the CERN openlab Technical Workshop was held at the CERN Council Chamber and brought together a broad cross-section of the CERN openlab community, including experts from industry, academia, and research organisations. The event highlighted the continued importance of public-private partnerships in advancing scientific research and addressing complex technological challenges.

Participants engaged in in-depth discussions on CERN openlab’s ongoing projects, exploring both technical progress and broader challenges at the interface of science and computing. The workshop also provided a valuable forum for reflecting on emerging trends, identifying future needs, and discussing opportunities for new and strengthened collaborations.

The programme covered a wide range of topics aligned with CERN openlab’s R&D activities, including heterogeneous computing platforms and infrastructures, artificial intelligence and HPC convergence, cloud and storage technologies, and applications for society and the environment. Industry partners also contributed through dedicated presentations, sharing perspectives and experiences from collaborative projects.

Beyond project updates, the workshop served as an important venue for fostering dialogue and collaboration, reinforcing CERN openlab’s role as a platform for co-development between industry and the research community and for shaping future directions in scientific computing.



Participants of the 2025 CERN openlab Technical Workshop photographed in front of building 500.



Outreach Events

Throughout 2025, CERN openlab actively pursued outreach opportunities through participation in major international conferences and events, supporting the dissemination of results and engagement with the wider scientific and technological communities. CERN openlab representatives contributed through presentations, keynotes, and panel discussions at events such as the EuroHPC Summit and EuroHPC User Days, Oracle CloudWorld, the InPEX Workshop, ISC High Performance, and the PASC Conference. These engagements helped strengthen CERN openlab's visibility, foster collaboration, and share insights on advanced computing, digital infrastructures, and emerging technologies with diverse international audiences.



Maria Girone, head of CERN openlab, presenting at the EuroHPC User Days in Copenhagen.



Maria Girone, head of CERN openlab, presenting at the ISC High-Performance conference in Hamburg.



Luca Atzori, CTO for Computing, presenting at the ISC High-Performance conference in Hamburg.



Maria Girone, head of CERN openlab, presenting at Oracle Cloud World tour in Zurich.



Maria Girone, head of CERN openlab, presenting at the EuroHPC Summit in Krakow.



Maria Girone, head of CERN openlab, presenting at InPEX in Kanagawa.

ESPPU Contribution

CERN openlab also contributed to broader community efforts in 2025, including inputs to the European Strategy for Particle Physics Update (ESPPU) and contributions to the proceedings of CHEP24, supporting strategic reflection and knowledge dissemination within the high-energy physics and computing communities.



Maria Girone, head of CERN openlab, presenting at the ESPPU meeting in Venice.

CHEP24 Contribution


Atzori, L., Girone, M., James, T., Mascetti, L., Nappi, A., Verder, K., Wulff, E. (2025). Industry-science R&D projects for the High-Luminosity LHC under CERN openlab. In EPJ Web of Conferences, 337, Article 01307. Presented at the 27th International Conference on Computing in High Energy and Nuclear Physics (CHEP 2024). https://www.epj-conferences.org/articles/epjconf/abs/2025/22/epjconf_chep2025_01307/epjconf_chep2025_01307.html

Digital Communication

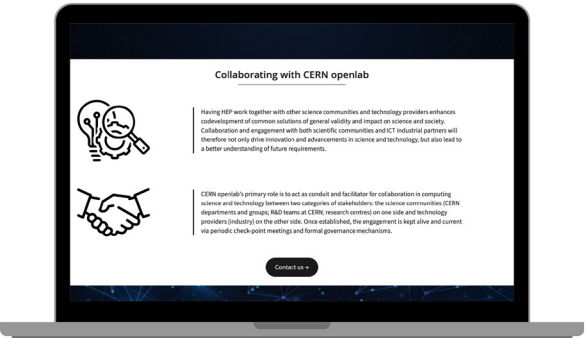
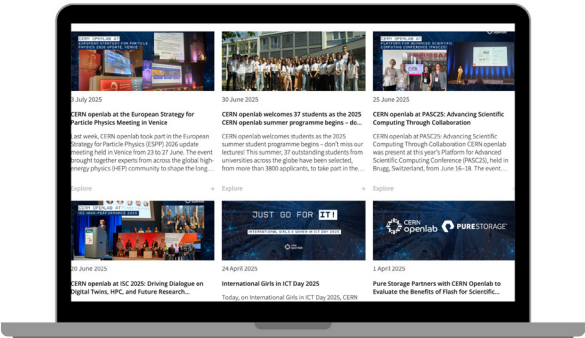
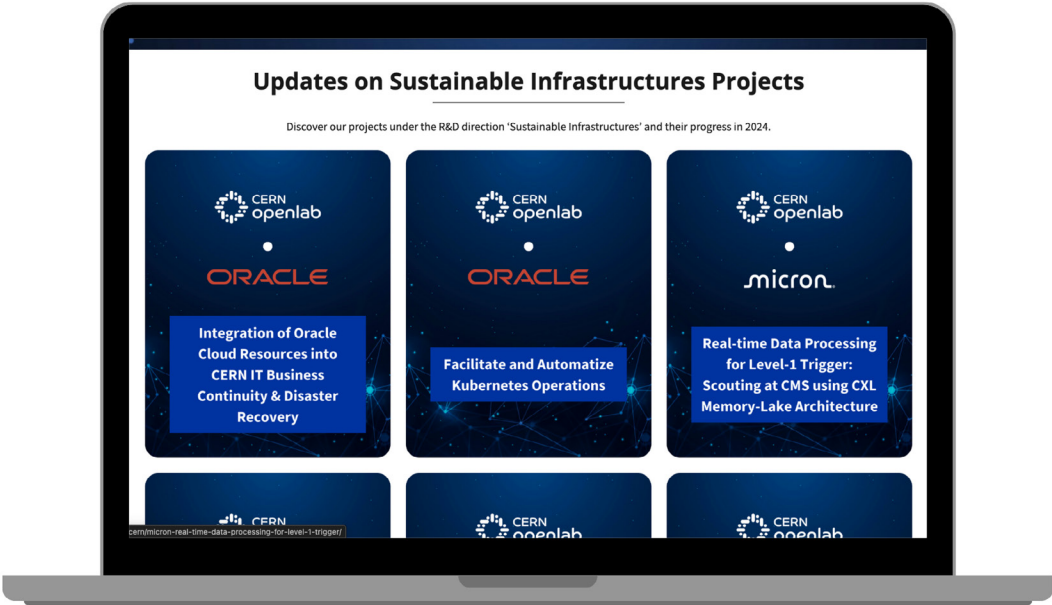
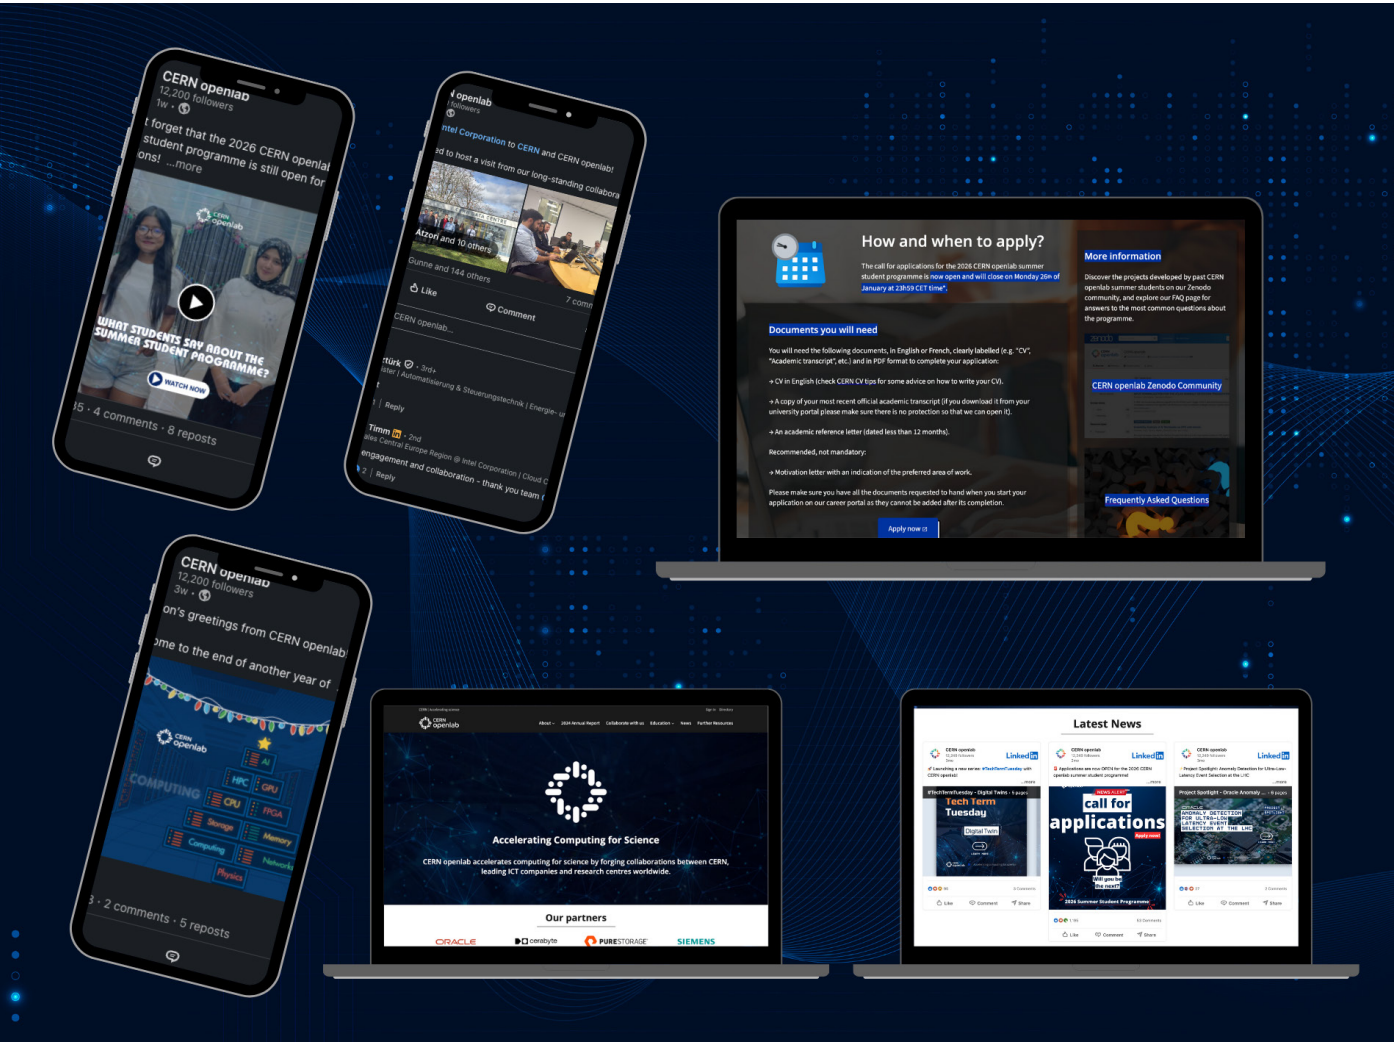
Digital communication plays a central role in CERN openlab’s outreach strategy, supporting visibility, transparency, and engagement with a broad scientific and technological audience. CERN openlab’s primary digital channels are its website and LinkedIn account, which together serve as key platforms for sharing news, project results, events, and opportunities. The website provides structured access to information on activities and collaborations, while LinkedIn enables timely updates and interaction with the wider research, industry, and innovation communities.

Website

In 2025, CERN openlab completed the migration of its website from Drupal to WordPress, in line with CERN’s broader transition to a new web platform. This milestone enabled the launch of a renewed [CERN openlab website](#) with a refreshed visual identity and improved structure, designed to better reflect the programme’s activities and priorities. The new layout allows visitors to easily explore CERN openlab projects, access the latest news and updates from the homepage, and consult a digital version of the annual report. Overall, the new website strengthens CERN openlab’s digital presence, improves accessibility to information, and supports more effective communication with partners, collaborators, and the wider community. Between February 2025 and January 2026, the CERN openlab website counted with more than 40,000 visits.



CERN openlab Website




New CERN openlab website ([openlab.cern](#)), showcasing key pages including the 2024 Annual Report (top), News (bottom left), and Collaborating with us (bottom right).

LinkedIn

LinkedIn has become a key communication channel for CERN openlab, supporting timely outreach and engagement with the wider scientific, technological, and innovation communities. The CERN openlab LinkedIn account was launched in July 2024 and saw significant growth throughout 2025, reaching more than 10,000 followers and generating over one million impressions. Through regular updates highlighting projects, events, collaborations, and opportunities, LinkedIn has strengthened CERN openlab’s digital visibility, enabled broader dissemination of activities, and fostered interaction with partners, researchers, and stakeholders worldwide.



CERN openlab
LinkedIn



In 2025, CERN openlab planned and implemented multiple targeted LinkedIn campaigns to increase visibility and engagement across its activities. These included dedicated campaigns for the CERN openlab summer student programme and the ‘Meet the Team’ series, highlighting both opportunities and the people behind CERN openlab. In addition, new recurring formats were introduced, such as ‘Project Spotlight’, which showcases CERN openlab projects through videos and slideshows, and ‘Tech Term Tuesday’, designed to explain computing concepts in a clear and accessible way. Together, these initiatives supported consistent storytelling, broadened outreach, and strengthened engagement with a diverse online audience.



Meet the Team

‘Meet the Team’ was a LinkedIn campaign that highlighted the people behind CERN openlab, showcasing their roles, expertise, and contributions to the programme and its collaborative projects




#TechTermTuesday

‘Tech Term Tuesday’ is a recurring LinkedIn campaign designed to demystify computing and digital technologies by presenting key technical terms in a clear, accessible format for a broad audience.





Project Spotlight

Project Spotlight is a LinkedIn campaign that features CERN openlab projects through short videos and visual content, highlighting their goals, progress, and technological impact in an accessible way.



2025 CERN openlab Summer Student Programme

The 2025 CERN openlab summer student programme was highlighted through a dedicated LinkedIn campaign, showcasing student projects and experiences.





Outreach Materials

CERN openlab produces a range of outreach materials to support the clear communication of its activities, results, and collaborations. These materials include news articles, videos featuring project outcomes and industry partners, promotional assets, and the annual report, all designed to present technical work in an accessible and engaging way. Together, they support knowledge dissemination, enhance visibility, and ensure that CERN openlab’s research and partnerships are effectively communicated to a wide range of audiences.


Videos with Industry

CERN openlab also produced a series of videos with industry partners, highlighting collaborative projects, shared expertise, and the impact of public-private partnerships in advancing scientific computing.



Annual Report

The CERN openlab Annual Report provides a comprehensive overview of activities, projects, and partnerships, serving as a key reference for communicating results and impact to the wider community.



Promotional Items

In 2025, CERN openlab developed a new range of promotional items (including T-shirts, tote bags, water bottles, mugs, and lanyards) to support outreach activities, increase visibility, and reinforce the project’s visual identity at events and within the community.



With thanks to

All partners who have collaborated with CERN openlab activities and everyone who has contributed to the content and production of this document.

Editors

Maria Girone
(Head of CERN openlab)

Mariana Velho
(CERN openlab communication, education and outreach manager)

Graphic Design & Layout

Mariana Velho
(CERN openlab communication, education and outreach manager)

Get in touch

 [linkedin.com/showcase/cernopenlab](https://www.linkedin.com/showcase/cernopenlab)

 openlab.cern

 openlab-communications@cern.ch



CERN openlab
LinkedIn



CERN openlab
Website



CERN openlab
Phase VIII Strategy



CERN openlab
2024 Annual Report

All images by Freepik, except:

Page 3, 4, 15, 16, 17, 18, 20, 22, 26, 32, 38, 39, 40, 45, 53, 54, 55, 56, 59, 60: CERN, CERN openlab members, CERN openlab communication office, CERN openlab research team members

Page 1, 7, 12, 29: Antonio Vivace for CERN openlab

Page 31: Cerabyte Ceramic-on-Glass Media (Image credit: Cerabyte)

Page 41: Alveolar segment model in TOPAS-nBio from Mechanistic model of radiotherapy-induced lung fibrosis using coupled 3D agent-based and Monte Carlo simulations

Page 44: Pasqal ORION quantum processor (Image credit: Pasqal)

Page 49: A pulsar (pink) can be seen at the center of the galaxy Messier 82 in this multi-wavelength portrait. The pulsar was discovered by NASA's NuSTAR which detected the pulsar's X-ray emission. (Image credit: NASA/JPL-Caltech)

Page 61: Maria presenting at EuroHPC Summit in Copenhagen (Image credit: Julie de Bellaing)

Page 61: Maria presenting at ISC High-Performance in Hamburg (Image credit: Ian Fisk)

Page 62: Maria presenting at EuroHPC User Days in Krakow (Image credit: Jakub Moscicki)

Page 62: Luca Atzori presenting at ISC High-Performance in Hamburg by (Image credit: Alexander Zoechbauer)

Page 62: Maria presenting at Oracle CloudWorld Tour in Zurich by (Image credit: Andrzej Nowick)

Page 62: Maria Girone, head of CERN openlab, presenting at InPex in Kanagawa (Image credit: Tommaso Boccali)

Page 62: Maria Girone, head of CERN openlab, presenting at the ESPPU meeting in Venice (Image credit: Sinéad Ryan)

DOI

<https://doi.org/10.5281/zenodo.18455533>

ISBN

978-92-9083-721-3 (Digital)
978-92-9083-720-6 (Printed)



©CERN 2026