CERN openlab

PHASE VIII (2024-2026)

ACCELERATING COMPUTING FOR SCIENCE

CERN openlab

CERN

Published by CERN

TABLE OF CONTENTS

Editor
Maria Girone
(Head of CERN openlab)

Graphic Design & Layout
Mariana Velho
(CERN IT Communication Officer)

Get in touch
openlab.cern
linkedin.com/showcase/cernopenlab
x.com/cernopenlab
facebook.com/cernopenlab
openlab-communications@cern.ch

ISBN
978-92-9083-651-3 (Digital)
978-92-9083-652-0 (Printed)

Published by CERN

A WORD FROM THE HEAD OF CERN OPENLAB

↑ **Maria Girone**, Head of CERN openlab

CERN openlab stands as a testament to over two decades of pioneering history and expertise. Established in 2001, it remains a ground-breaking **public-private partnership**, forging collaborations between leading ICT companies and research centres worldwide, uniting them with the forefront of scientific innovation at CERN. These partnerships fuel CERN researchers with invaluable opportunities and resources to push the boundaries of computing, essential for tackling the unprecedented challenges presented by monumental projects like the High-Luminosity Large Hadron Collider (HL-LHC) or the Square Kilometre Array (SKA).

For more than twenty years, CERN openlab has been the bedrock of partnership, providing a structured and collaborative framework for industry and research organisations to engage with CERN researchers. Its framework not only fosters collaboration but also facilitates industry participation and investment.

↓ CERN openlab team members

Collaboration is the core of CERN openlab, where industry and scientific researchers co-develop innovative solutions. Together, we embark on four primary missions: establishing strategic industry collaborations, fuelling technological innovation, exposing technology to researchers, and nurturing knowledge and growth in young STEM researchers.

Now embarking on its eighth phase, with 23 years of transformative history, CERN openlab embraces a new era of multidisciplinary collaboration, with fresh insights and activities. As the world evolves and new challenges emerge, our dedication to accelerate scientific computing, propelling technologies beyond HEP, and fostering positive societal impacts on a global scale remains unwavering.

Maria Girone
Head of CERN openlab

CERN OPENLAB MISSION

Since its inception, CERN openlab has fostered the development of big data scientific research through **four primary missions**:

1 Establishing strategic industry collaborations

CERN openlab projects provide an ideal **incubator for collaborations**, where they can be formed and longer-term partnerships can be built. They act as the first step in the establishment of strategic collaborations between CERN and other organisations interested in investing in the future of science and technology.

2 Fuelling technological innovation

CERN openlab serves as an **incubator for new technologies**. It forms a dynamic hub where CERN and its partners collaboratively push the boundaries of ICT technology. This cooperative synergy propels the co-development of new ideas and innovative solutions.

3 Exposing technology to researchers

CERN openlab provides access to new technologies available on the market to its members and the HEP community, **supporting critical tasks** of evaluation, adaptation, and benchmarking.

4 Nurturing knowledge and growth in young STEM researchers

CERN openlab plays a crucial role in **training the experts of tomorrow**. The CERN openlab summer student programme, supported through industry contributions, trains students through real-world, concrete, multi-disciplinary projects. Through various programmes and workshops, it equips the next generation of researchers with essential skills required to navigate the complex landscape of modern computing technologies.

To address scientific challenges at the exascale level, CERN openlab has identified two main R&D directions: "**Sustainable Infrastructures**" and "**Emerging Technologies**". In sustainable infrastructures, it collaborates on energy-efficient computing platforms, while in emerging technologies, it explores innovations like long-term digital storage materials, scientific digital twins, and quantum computing.

This document explains how CERN openlab has developed, how it intends to operate, which communities are targeted, what objectives it sets for itself and its collaborations, and how it proposes to continue delivering **impactful results and sustainable collaborations** for the years to come.

CERN OPENLAB COLLABORATION MODEL

CERN openlab operates within structured **three-year phase cycles** designed to systematically assess technological evolution, anticipate future needs, and delineate overarching thematic priorities. This approach ensures the maintenance of a relevant and current research programme, fostering effective collaborations and innovative advancements.

Upon joining CERN openlab, members gain access to a unique ecosystem characterized by **unparalleled computing challenges, ground-breaking scientific endeavours, and pioneering minds**. This environment serves as a crucible for the development and demonstration of emerging technologies, providing a platform for industry leaders to showcase their potential and validate solutions through realistic, demanding use cases. This process often leads to tangible enhancements in product features and capabilities.

In addition to impactful technical progress, CERN openlab members' have a valuable opportunity to boost their global reputation. Affiliation with a renowned institution like CERN elevates participants' visibility and stature within broader scientific and technological communities. CERN openlab represents the **convergence of cutting-edge science, collaborative exploration of disruptive technologies, and extensive visibility**, resulting in tangible strategic benefits for its members.

EVOLUTION OF THE COLLABORATION MODEL

INCEPTION

PHASES I-IV

CERN openlab was founded in 2001 at a time of great technological upheaval due to the transition from mainframe to commodity computing. Initially, large industrial partners were attracted to participate and invest in the design and construction of the computing infrastructures for the LHC. A number of top industry leaders joined CERN openlab, some of which are still actively involved. With the help of companies like Intel, IBM, and HP, CERN openlab navigated the challenges of transitioning the CERN IT data centre to modern x86 64-bit processors, multi-core devices, high-performance networking, and petabyte storage. CERN openlab's collaboration with Oracle helped to establish high availability and redundancy for the critical online databases needed for detector operations. At the time, CERN was doing something nobody else was doing: defining the new **worldwide reference for big data scientific research**.

CONSOLIDATION

PHASES V-VII

The second decade of CERN openlab coincided with LHC operation and a need for higher stability and reliability. At the same time, the technology market became more dynamic and fragmented, with a larger number of smaller companies and even start-ups showcasing innovative technologies. The CERN openlab operating model evolved to focus on shorter development/evaluation cycles and more targeted technology interests. We expanded the programme for the benefit of the accelerator and detector operations and defined a new model to collaborate with universities and research centres sharing challenges at the data and computer science level with a broader multi-disciplinary view. **The number of partners grew towards a broader portfolio of smaller agile projects**.

THE NEXT PHASES

PHASE VIII AND BEYOND

CERN, along with the High Energy Physics (HEP) community, is gearing up for the HL-LHC program and the future of particle physics. Building upon the achievements of the initial decade of LHC operations, marked by swift advancements in computing architectures, infrastructures, algorithms, and disruptive technologies, CERN openlab is poised to harness these changes. With its extensive network spanning industry and academia, CERN openlab is uniquely positioned to validate and integrate emerging technology capabilities while facilitating access to novel infrastructures.

To address the impending challenges, **CERN openlab is reviewing its operational framework**, notably restructuring industry membership into two tiers: Associate and Partner (refer to membership levels and governance sections).

CERN openlab will build on its consolidated agile mechanism that facilitates the definition of practical projects, to remain an efficient vehicle for innovation. Projects typically span between 1 to 3 years, an optimal duration to yield actionable outcomes aligned with the latest technological advancements. This framework represents an **ideal mechanism to nurture initial collaborations that can evolve into long-term programmes, but also to promote new partnerships and cultivate diverse networking avenues**.

In the coming years, CERN openlab will introduce additional tools to facilitate sustained, long-term collaborations, serving as a platform not only for partnership preparation but also as a "training ground" for IT professionals seeking to expand their skill sets into project management and broader coordination roles.

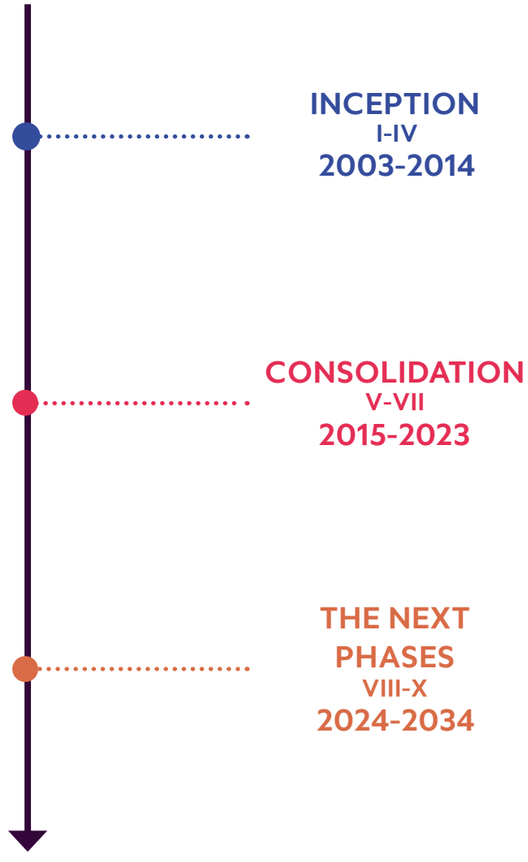

MOTIVATIONS, OBJECTIVES, & EXPECTED IMPACT FOR MEMBERS

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.

CERN openlab **objectives** are to pioneer sustainable and innovative computing solutions, harness AI and heterogeneous computing for environmental benefits, and foster collaboration and technology transfer between industry and the scientific community. Through collaboration with diverse stakeholders, including the HEP community, other scientific disciplines, and technology providers, CERN openlab fosters co-development of solutions and co-design of infrastructure. 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 **R&D directions of "Sustainable Infrastructures" and "Emerging Technologies"**.

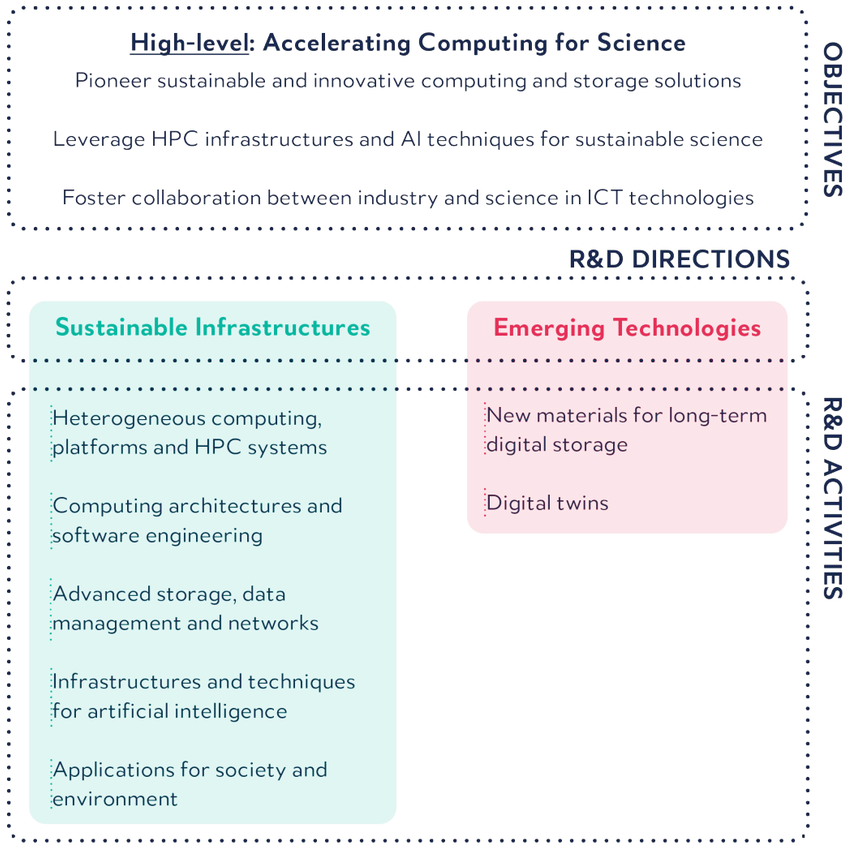

OUR R&D ACTIVITIES
SUSTAINABLE INFRASTRUCTURES

Heterogeneous computing, platforms and HPC systems

CERN openlab contributes to this area of innovation via collaboration with resource providers for access to test resources on cloud and HPC infrastructures. The main focus is on understanding how to access and integrate external resources into the computing workflows employed by the CERN community (e.g., IT, experiments). This can be done via direct access to cloud environments or pilot collaborations with cloud providers and HPC sites. This activity includes optimisation of AI workflows on large-scale HPC systems.

Impact

The expected impact is the **co-development of adapted access models to HPC and cloud infrastructures, integrated and optimised workflows, including AI, and costing models**.

Computing architectures and software engineering

CERN openlab contributes to this area of innovation via collaboration with resource providers for the provision of hardware and software components, dedicated specialised expertise from the technology partners, funds for engineers and developers, as well as training and education opportunities. The main focus is on assessment, benchmarking and validation of accelerated architectures (e.g. GPUs, FPGAs) and new processors (e.g. RISC V) following user requirements.

Impact

The expected impact is to provide the HEP community with **innovative technology and key expertise on ongoing and future research programmes**.

Advanced storage, data management and networks

CERN openlab contributes to this area of innovation via collaborations with technology providers and systems integrators. This includes, for example, the evaluation of new storage media, the co-development of specific functionality for multi-disciplinary applications, or the definition and implementation of data workflows for HEP in the context of the emerging data analysis facilities concept.

Impact

The expected impact is to provide the HEP community with **storage and data management tools that can effectively support cutting-edge research programmes**.

Infrastructures and techniques for artificial intelligence

CERN openlab contributes to this area of innovation by fostering collaboration with technology providers and research institutions developing state-of-the-art platforms, services, and methodologies. This includes access to software and expertise, as well as large-scale testbeds for co-creating new AI models and optimization workflows. The areas of work include distributed AI optimization, generative AI, foundation models for physics, as well as optimal deployment of AI-based algorithms on modern computing architectures, and benchmarking of these architectures on AI workloads. Work is also ongoing in the realm of real-time AI inference (on edge or accelerator devices) as part of the detector data acquisition, fast data selection, and accelerator control.

Impact

The expected impact is to **enable the science community to access and leverage AI resources and skills**.

Applications for society and environment

Energy-efficient computing and AI for society are of great interest to technology providers, data centres, and software developers. CERN openlab will contribute both in the establishment of dedicated investigations in green computing, life and environmental science, and in raising awareness through explicit requirements in the definition of development projects.

Impact

CERN openlab will support the migration to **energy-efficient architectures** and develop tools for **sustainable data centres**. Furthermore, it will keep its work with other sciences and contribute to research with a societal impact.

↓ CERN tape storage system

↓ AI generated image of a woman working on an accelerator digital twin

EMERGING TECHNOLOGIES

New materials for long-term digital storage

On the time-scale of HL-LHC, it is expected that tape will remain the archive medium of choice for CERN experiments. However, other archival technologies are in the R&D stage and should be investigated for possible future exploitation. This topic evaluates the characteristics and likely evolution of **new materials** as archival storage mediums to inform CERN's long-term plans for data archival.

Impact

CERN openlab will foster investigations of **innovative, long-term, and durable storage solutions (DNA, ceramics, silicon).**

Digital twins

Digital Twin technology is gaining traction across a diverse range of fields. It involves creating a dynamic digital counterpart of a physical entity, continuously updated to mirror real-world changes. CERN openlab focuses on Digital Twin use cases for scientific applications. Our involvement in the interTwin project, funded by the European Commission, focuses on developing a versatile digital twin engine prototype. This initiative spans physics, including particle physics, radio astronomy, and gravitational waves, as well as environmental modelling for early warnings using satellite imagery. Another direction of exploration includes CAD-based digital twins to enhance the operation of the Large Hadron Collider (LHC) and experiments. This not only aids in maintenance but it would also allow to simulate detector assembly scenarios. The potential applications are vast, offering simulation of underground interventions in a virtual space, resulting in **cost savings, time efficiency, and reduced exposure to radiation**.

Impact

CERN openlab will leverage the expertise developed in digital twin applications to offer **versatile solutions, from scientific applications to enhancing LHC operations, enabling real-time modelling of real-world scenarios, predictions, and time-cost optimization**.

IMPLEMENTATION MODEL

The CERN openlab **implementation model relies on two main approaches**:

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

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

Leveraging these approaches, CERN openlab actively pursues collaborations with industry and academia, establishing a foundation of trust for enduring partnerships. This approach not only maximizes technical impact but also fully harnesses the potential of innovative technologies, thereby amplifying our core strengths.

STAKEHOLDERS

Having HEP work together with other science communities and technology providers enhances co-development of common solutions of general validity and impact on science and society. Collaboration and engagement with both scientific communities and ICT industrial partners will therefore not only **drive innovation and advancements in science and technology**, but also lead to a better understanding of future requirements.

CERN openlab's primary role is to **act as conduit and facilitator for collaboration in computing science and technology between two categories of stakeholders**: the science communities (CERN departments and groups; R&D teams at CERN; research centres) on one side and technology providers (industry) on the other side. Once established, **the engagement is kept alive and current via periodic check-point meetings and formal governance mechanisms** (more information on the governance section).

MEMBERSHIP LEVELS

For CERN openlab Phase VIII (2024-2026), a **simplified membership structure is being implemented** with two levels: strategic members (Partners) and associate members (Associates).

	PARTNERS	ASSOCIATES
In-kind contributions & annual funded effort	≥300 000CHF	<300 000CHF
Annual membership	30 000CHF	30 000CHF
Scope of collaboration	Longer-term co-development programmes	Shorter tactical projects

Projects are required to **include overhead costs for activities such as system administration, dedicated communication and support activities**. The appropriate level of overhead to be charged depends on the type of project and will be explicitly negotiated, but it is estimated to be of the order of 10% of the total project cost.

Although the main distinction between Partners and Associates is in the scope of the collaboration, Partners have more benefits than Associates. **Participation in the CERN openlab Collaboration Board is reserved to Partners**.

Research centres with complementary expertise, aligned scientific goals, and innovation objectives may become research members of CERN openlab with or without direct financial contributions.

OUR TEAM & GOVERNANCE

CERN openlab operates as a dynamic bridge between innovation and practical application, functioning both as a broker and an incubator. Its ongoing success and sustainability are underpinned by contributions from its member organisation, including significant direct support from CERN, particularly in areas of openlab management and coordination efforts. The CERN openlab Governance relies on a **small central coordination team and support from specific IT functions**. The coordination team works in close collaboration with CERN IT teams, including IT Collaborations and Partnerships, the Project Management Office, Engagement, Technical Delivery, and Communication. Regular meetings occur between the Head of CERN openlab, the CTO office, administrative, communication, finance support, and the projects' coordinators. Highlighting our commitment to collaboration and transparency, we convene an annual Collaboration Board meeting. This event brings together industry representatives at the Partner level and coincides with the CERN openlab Technical Workshop. It serves as a platform to review and celebrate the progress and achievements of our projects, fostering a shared sense of accomplishment and direction among all stakeholders.

Head of CERN openlab

Coordinates the overall programme and manages the negotiation of contracts between CERN and members with the assistance of the CERN Legal Services.

Maria Girone
Head of CERN openlab

CTO office

Contributions from experts in CERN IT technical groups. Assesses technology for projects. It functions as a central hub for proposals evaluation and project coordination and maintains close ties with CERN technical groups and other CERN departments.

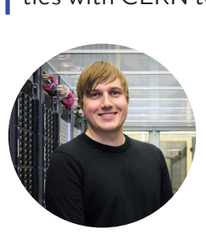

Thomas Owen James
CTO for AI and Edge Devices

Antonio Nappi
CTO for Platforms and Workflows

Luca Mascetti
CTO for Storage

Luca Atzori
CTO for Computing

Killian Verder
CTO Office Administrative

Communication office

Manages CERN openlab's communication needs, especially agreements with industrial partners about communication, including organising the summer student programme and events.

Mariana Velho
Communication, Education & Outreach Manager

Aministrative & Financial office

Handles administrative tasks.

Valentina Clavel
Finance Manager

TRAINING & EDUCATION

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

Apart from that, the public has open access to CERN openlab **lectures that cover a wide range of computing topics, from AI to exascale computing and quantum technologies**. For community development, CERN openlab offers regular specialised technical training to members of the scientific community. Recent examples include NVIDIA GPU programming with CUDA hackathons, Intel software tools, hands-on training for unified programming and AI/ML models, and containerisation and virtualisation training including Kubernetes.

When preparing the future of scientific computing, it is vital to **ensure that the computing specialists of the future have the right skills to enable them to fully capitalise on new, innovative technologies**. Through projects, lectures, and workshops, computer scientists are being equipped with indispensable knowledge that inspires scientific advancement and fuels innovation.

↓ Students from past CERN openlab summer student programmes

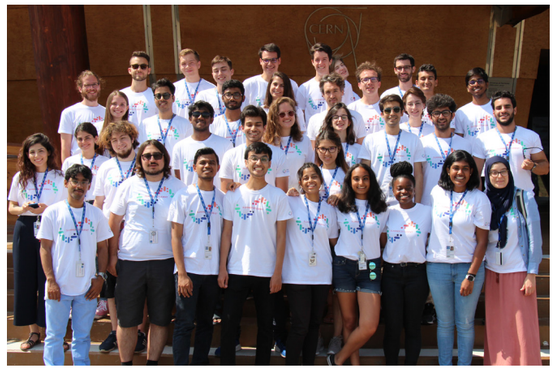

WITH THANKS TO

all CERN openlab partners and collaborators and everyone who has contributed to the content and production of this document.

ALL IMAGES PROVIDED BY THE CERN AUDIO-VISUAL PRODUCTION SERVICE EXCEPT:

woman working on a detector digital twin (page 9) - by Pixlr's AI Image Generator.

GET IN TOUCH

openlab.cern

facebook.com/cernopenlab

x.com/cernopenlab

openlab-communications@cern.ch

linkedin.com/showcase/cernopenlab

WEBSITE FACEBOOK TWITTER LINKEDIN

READ MORE ABOUT THE CERN OPENLAB PROJECTS ON THE DIGITAL CERN OPENLAB 2023 ANNUAL REPORT

Editor
Maria Girone
(Head of CERN openlab)

Openlab Contributors
Luca Atzori
Marina Banjac
Matteo Bunino
Alberto Di Meglio
Kristina Gunne
Thomas Owen James
Luca Mascetti
Antonio Nappi
David Southwick
Mariana Velho
Killian Verder
Eric Wulff
Alexander Zoechbauer

Graphic Design & Layout
Mariana Velho
(CERN IT Communication Officer)

ISBN
978-92-9083-651-3 (Digital)
978-92-9083-652-0 (Printed)

Published by CERN