




Implementation and Resource Plan
Artificial Intelligence at CERN
A Core Capability

2026-2030



Table of Contents

1. Executive Summary	3
1.1 Strategic Intent	3
1.2 Priorities and Timelines	4
1.3 Resource Estimates	5
2. Implementation Model and Principles	7
2.1 Relationship to the AI Strategy	7
2.2 The Implementation model	7
2.3 Transversal Principles for Implementation.....	8
3. Strategic Objectives and Expected Outcomes.....	10
3.1 Pillar 1: AI as a Scientific Enabler	10
3.2 Pillar 2: AI as an Operational and Organisational Enabler	12
3.3 Pillar 3: AI as an Impact and Collaboration Enabler.....	13
3.4 Lighthouse Initiatives	15
4. Phased Roadmap	18
4.1 Phased Timeline	18
5. Governance.....	21
6. Resource Plan	22
7. Implementation and Operations.....	23
Lighthouse Initiative Factsheets	24

1. Executive Summary

1.1 Strategic Intent

CERN's mission is to advance understanding of the fundamental constituents and laws of the Universe, operating the world's largest particle accelerator complex and serving as host laboratory for the global high-energy physics community. Delivering the scientific ambition of the LHC, its high-luminosity upgrade, and the broader research programme increasingly depends on extracting knowledge from data of a scale and complexity that grow faster than computing and human effort alone can absorb. Artificial Intelligence has already become part of how CERN does science and runs its infrastructure, from data analysis and simulation to accelerator operations and possibly in the future accelerators and detector design, with concrete projects or R&D initiatives already underway across the CERN sectors and departments, and the experiments. The recent emergence of foundation models and agentic AI has widened these possibilities, from scientific discovery and detector design to knowledge management and decision support and is a principal reason for setting out a coordinated approach now.

CERN's role in AI is twofold: to apply AI in the service of its own scientific and operational mission as host laboratory, and to contribute, as it has historically done for computing and software, to the shared AI capabilities of the worldwide high-energy physics community. This Implementation Plan is designed to support this scientific purpose. CERN pursues this in complementarity with the worldwide HEP community and other AI-for-science initiatives, engaging through established community fora and sharing open, reusable methods and infrastructure rather than duplicating efforts.


Over the past decade, CERN has developed significant expertise in AI through numerous initiatives across departments, and services. These efforts have demonstrated substantial benefits, but have often evolved independently, leading to fragmented developments and duplicated investments. At the same time, the emergence of large-scale AI models and new computing paradigms requires coordinated planning, shared infrastructure, and long-term strategic investments.

CERN faces a convergence of structural challenges — in data volume, operational complexity, sustainability and resource constraints, and scientific ambition — that cannot be resolved just by scaling existing approaches. The CERN AI Strategy 2026–2030¹, endorsed by CERN Management in March 2026, establishes Artificial Intelligence as a core organisational capability, complementing CERN's established strengths in computing, detectors, and engineering. This Implementation Plan translates the strategic commitment into a phased programme of actions, with a defined resource envelope, a governance model, and a set of measurable outcomes against which progress will be reported annually. Beyond deploying existing AI technologies, the plan's ambition is to establish CERN as a leading laboratory for AI-enabled physics research, with key objectives including foundation models for scientific data, agentic systems supporting the design and operation of detectors, accelerators, and computing services, and intelligent assistants that enhance scientific productivity and knowledge access. Realising this will require CERN's computing and data infrastructures to evolve to support

¹ <https://repository.cern/records/zw9za-jqv67>

modern AI workloads, including large-scale training, distributed inference, model serving, and data governance.

The plan also identifies a small set of lighthouse initiatives: flagship, cross-cutting efforts that each draw on several of the plan’s objectives and serve to illustrate what they could make possible. They are presented in two tiers. A first group of short-medium term initiatives is the more mature group, expected to deliver results within the plan’s technology, timeline, and funding, and includes work towards a common AI infrastructure and platform, information discovery and retrieval, real-time AI at the edge for data acquisition, fast simulation, AI support for accelerator operations, and physics foundation models. A second group of horizon initiatives is deliberately more exploratory and longer-term, likely to mature mostly beyond 2030 and dependent on resources and partnerships not initially committed in this plan. The set is indicative rather than exhaustive and will be reviewed and adjusted following achievements, progress, and technology evolution.



1.2 Priorities and Timelines

The implementation is structured in four successive phases with prioritised actions and designed to deliver progressively tangible outcomes. Figure 1 shows the phased timeline and the principal milestones associated with each phase. The schedule shown is indicative and will be

updated as the activities progress; the main text is concerned only with the high-level objectives for the 2026 to 2030 mandate.




Figure 1: Prioritised phased timeline 2026-2030

These phases build on AI work already ongoing across CERN for several years; the plan's aim is to accelerate that progress, coordinate it within a shared framework, and improve the sharing of know-how across the Organisation. Throughout, the plan takes as established the machine-learning methods already in long-standing use across high-energy physics, and concentrates on the more recent generative and agentic AI, where the principal new requirements and opportunities lie. Phase 1 (Q3 2026–Q1 2027) establishes the governance bodies, shared data foundations and AI platform infrastructure. Phase 2 (Q2–Q4 2027) transitions these foundations into operational capabilities, with data governance, the responsible AI policy, and advanced training tracks becoming fully active. Phase 3 (2028) delivers the first measurable scientific and operational results, including initial scientific outcomes from RCS and first pilots from ATS. Phase 4 (2029–2030) achieves full HL-LHC readiness, with production-grade AI systems operating across accelerator, research, and operational domains, and sustained governance and review cycles maintaining alignment with CERN's mission through 2030.

The phased approach delivers tangible results early, establishes natural checkpoints for annual portfolio reviews, and makes this plan a living document allowing maximum flexibility for adjustments as the programme matures, technology evolves, and lessons from early phases are incorporated. Outcomes that depend on adoption by the experiment collaborations are pursued jointly with them, always seeking community agreement and support.

1.3 Resource Estimates

The resource plan accompanying this document provides an estimate of the resources already committed to ongoing activities and the additional cost of delivering the new identified

objectives over 2026–2030. Of the total estimated cost, a significant portion is met by existing internal and external resources already committed to ongoing AI pilots and investigations. Existing resources remain under the management of the respective sectors and departments and are presented as an indicative estimate rather than a commitment.

AI research, development, and deployment have been embedded in CERN’s scientific and technical activities for decades, and a substantial fraction of the expertise, personnel, infrastructure, and external collaborations relevant to this programme already exists within sectors, departments, experiments, and services.

Existing resources imply no transfer of personnel, budgets, reporting lines, or responsibilities to a central AI structure, and ownership, prioritisation, and management of these activities remain with the respective sectors, departments, experiments, and services.

The additional investment from CERN is intended to complement, not replace, existing efforts. Their primary purpose is to establish the common capabilities that individual groups cannot efficiently develop in isolation, including shared AI infrastructure, common data and model services, training programmes, governance frameworks, benchmarking activities, and cross-organisational strategic initiatives. By providing these common foundations, the programme aims to accelerate ongoing developments, reduce duplication of effort, improve knowledge sharing, and enable new collaborations across CERN while preserving the autonomy and scientific priorities of the participating communities.

2. Implementation Model and Principles

2.1 Relationship to the AI Strategy

This Implementation Plan operationalises the CERN AI Strategy 2026–2030. Where the Strategy defines the vision, principles, and strategic objectives for AI as a core organisational capability, this document translates them into sequenced actions, a phased timeline, a governance model, and an estimated resource plan. It does not centralise scientific direction nor prescribe specific technologies. Execution responsibility remains with CERN departments, operating within a shared framework that ensures coherence, avoids duplication, and enables progress to be reported and reviewed at the level of CERN Management. Decision-making remains at the lowest appropriate level; only matters with significant cross-organisational, resource, or institutional implications are escalated to the C4.



Figure 2: The AI Strategy pillars and objectives

2.2 The Implementation model

The three pillars defined in the AI Strategy — AI as a Scientific Enabler, AI as an Operational and Organisational Enabler, and AI as an Impact and Collaboration Enabler — define the scope of the implementation plan. Collaboration with the experiment collaborations and the wider HEP community is addressed under Pillar 1; Pillar 2 concerns efficiency and reliability gains; Pillar 3 concerns CERN's broader scientific, societal, and industrial impact. Within each pillar, a set of specific objectives establishes the outcomes CERN is committing to deliver and the indicators

against which progress will be measured. These objectives are realised through two complementary types of action:

1. cross-organisational structuring actions that build shared, CERN-wide foundations regardless of sector, and
2. sector-specific actions and joint development actions that address the needs and priorities of each domain, seeking synergies whenever appropriate.

Both types of action are sequenced into a four-phase timeline designed to deliver progressively tangible results from 2026 through to HL-LHC readiness in 2030. A governance and portfolio review mechanism, operating under the principle of freedom in a framework established in the AI Strategy, ensures that strategic coherence is maintained and that progress, risks, and value creation are reported to CERN Management on an annual basis. The resource plan follows this structure directly, estimating financial and human effort action by action and phase by phase rather than from a fixed top-down budget.

Scientific AI applications developed within the experiment collaborations or research programmes remain subject to the established validation, review, and approval procedures of the respective scientific communities. The role of this Implementation Plan is not to introduce additional scientific review processes, but to facilitate the sharing of best practices, tools, and experience where appropriate.

Each pillar is further elaborated in Section 3, where objectives, expected outcomes, and impact indicators are set out. The extent to which deliverables can be fully realised depends on available resources and the continuous evolution of technology, external partnerships, and the international AI landscape. The implementation plan will accordingly be reviewed and updated on an annual basis to maintain its relevance and alignment with CERN's overall strategic objectives².

2.3 Transversal Principles for Implementation

A set of overarching transversal principles applies to all pillars, objectives, and actions, and is derived directly from the AI Strategy. These principles serve as the practical criteria against which implementation decisions on prioritisation, technology selection, resource allocation, and governance are assessed by the AI Programme Committee, the C4 and the CERN management (see section 5 for the Governance model).

Table 1: Implementation principles and criteria

Principle	Implementation criteria
Scientific Mission First	Implementation is driven first by the needs of the HL-LHC and the HEP scientific programme. CERN acts in its roles as host laboratory and as a contributor to the worldwide HEP community, respecting the autonomy of the experiment collaborations and sectors.

² <https://home.cern/news/opinion/cern/building-culture-success>

AI as a Core Capability	• AI driven by scientific questions and experimental constraints • Part of the scientific work required to perform research • Sustain the development of new scientific methods
Innovation and Long-Term Vision	• CERN and the HEP programmes span decades • Balance short-term optimisation and long-term breakthroughs
Measurable Value Creation	Every initiative must demonstrate one or more of: • Scientific impact • Operational efficiency gain • Sustainability benefit • Organisational simplification
Responsible and Trustworthy AI	• Mandatory validation standards for critical operations or systems • Human oversight in safety-critical systems • Clear accountability structures
Openness and Collaboration	• Open Science is part of the CERN way of working • CERN commits to collaborations that produce return on investment for its (Associate)Member States and advance social and economic value • CERN also commits to seeking internal synergies and opportunities for co-development across departments
Selective Internal Development	• Adopt mature external AI where appropriate • Develop internally only where CERN has unique differentiating value or it's justified by TCO
Sustainability by Design	• Energy-aware compute allocation • Lifecycle cost evaluation • Avoid technical debt accumulation • Adopt frugal AI principles across development and operations

3. Strategic Objectives and Expected Outcomes

Each pillar of the AI Strategy addresses a distinct but interconnected dimension of CERN's mission: the advancement of scientific discovery, the modernisation of operations and the organisation, and the extension of CERN's impact and leadership, contributing to and helping advance the AI methods, tools, and infrastructure of the worldwide HEP community, and helping to shape European AI-for-science initiatives, the Member States, and industry. As with initiatives CERN has led or helped develop, such as ROOT, CVMFS, RUCIO, INSPIRE, and Indico, solutions developed to address AI-related scientific and operational challenges have the potential to become widely adopted, within and beyond HEP. The objectives set out below are expressed as expected outcomes, each with an associated indicator and the phase by which the outcome is expected to be achieved. The detailed action plans, deliverables by phase, and full impact factor descriptions supporting each objective are described in the detailed Sector and Department Workplans.

This section is organised in two layers. For each pillar, a set of objectives expresses the specific outcomes CERN commits to, each with an indicator and a target phase. These are followed by a small number of lighthouse initiatives: flagship, cross-cutting efforts that each draw on several objectives and are the most visible demonstration of the plan. The objectives define what CERN will achieve; the initiatives illustrate what that achievement makes possible.

3.1 Pillar 1: AI as a Scientific Enabler

The HL-LHC and future collider facilities will place increasing demands on simulation, reconstruction, analysis, accelerator operations, and computing infrastructures. AI techniques offer powerful new approaches to addressing these challenges, enabling more efficient use of resources, improved data processing performance, and new ways of extracting information from complex datasets. While alternative solutions would remain possible, the effective adoption of AI has the potential to significantly enhance the scientific reach of the HL-LHC programme and to inform the design and feasibility assessment of future facilities such as the FCC. At the same time, AI opens qualitatively new modes of scientific inquiry, allowing exploration of parameter spaces and complex correlations that are inaccessible to conventional methods (for example, designing new detector concepts or discovering new materials for magnets) and, as a longer-term and more exploratory prospect, the possibility of helping to generate new physics hypotheses and theoretical models. CERN's commitment under this pillar is to institutionalise AI as a scientific capability complementing computing, detector R&D, and theoretical modelling with shared infrastructure, validated methods, and measurable outcomes. This work is carried out in close collaboration with the LHC experiment collaborations and the wider high-energy physics community, developing shared, reusable methods, datasets, and infrastructure rather than parallel or duplicated capabilities. In parallel, at every step, technical progress in AI will be assessed against CERN needs, so that always the optimal and latest-state technology is considered and used if it exists and co-developed by CERN in collaboration with the international community if it doesn't, always looking at long-term breakthrough potential beyond short-term optimisation.

Objective 1.1: Maximise the physics reach of the HL-LHC and enable research goals at current and future facilities

Building on the experiments' established practices in open data, common challenges, and shared detector models, CERN, with the experiment collaborations, will consolidate and extend common methodologies, datasets, cross-experiment benchmarks, and AI-driven data processing chains that increase the effective physics reach of the HL-LHC, with AI methods integrated into early design phases of future facilities by the end of the programme.

Key indicator: Measurable improvement in trigger efficiency, reduction in computational cost per event for processing and simulation, development of new AI-driven physics workflows and physics-driven generative AI models.

Expected by: Phase 3 (initial results), Phase 4 (full integration).

Objective 1.2: Accelerate data processing, improve signal extraction, optimise accelerator operations, and assist the design of next-generation devices

CERN will deploy pilot digital twin frameworks and shared surrogate modelling environments for selected accelerator systems, scaling to continuous model training on HPC resources and integration into accelerator operations by the end of the programme.

Key indicator: Demonstrated reduction in design iteration time; measurable performance gains in accelerator optimisation pilots.

Expected by: Phase 2 (pilots), Phase 4 (operational deployment).

Objective 1.3: Establish shared infrastructure, frameworks, and long-term support so AI is a core, scalable scientific capability at CERN

CERN will deploy a centralised AI development environment with a model registry, versioning system, and federated data access layer in Phase 2, scaling to full production across on-premises and distributed HEP and EU infrastructures by Phase 4.

Key indicator: Seamless execution of AI workflows as part of production activities on all necessary platforms.

Expected by: Phase 2 (pilot production), Phase 4 (full production).

Objective 1.4: Deploy validated, reproducible, and interpretable AI methods aligned with CERN's standards of openness and rigour

CERN will extend open model repositories with versioning and benchmark support and develop large-scale AI foundation models validated across the HEP community, with tools for reliable uncertainty quantification progressively available in production from 2027.

Key indicator: Percentage of production AI models documented and registered; availability of uncertainty quantification tools in active scientific workflows.

Expected by: Phase 3.

3.2 Pillar 2: AI as an Operational and Organisational Enabler

CERN's operational environment is growing in complexity, and many of its processes remain manual and labour-intensive. Accelerator systems and experiments, site infrastructure, and administrative functions offer scope for efficiency and reliability gains where AI can extend conventional automation, for example in predictive maintenance, real-time monitoring, process automation, and knowledge management. Realising these gains requires not only deploying AI tools but establishing the shared platforms, governance standards, and human oversight mechanisms that allow AI to be operated reliably and responsibly at scale. CERN's commitment under this pillar is to deploy AI as a measurable driver of operational efficiency and organisational simplification, without compromising safety, reliability, or institutional trust. A balance between industry-proven solution and potential for innovation and impact based on the selective development principle will guide the technology choices.

Objective 2.1: Use predictive and AI-assisted operations to increase availability, reduce downtime, and manage growing system complexity safely

CERN will inventory predictive maintenance opportunities and launch operational AI pilots in Phase 1, scaling to AI-enabled monitoring processes, dashboards, and quantified reductions in downtime and energy consumption by Phase 4.

Key indicator: Measurable increase in system availability and reduction in energy consumption attributable to AI-enabled operations; reduced pressure on operations personnel.

Expected by: Phase 2 (pilots at scale), Phase 4 (full production).

Objective 2.2: Implement shared, interoperable, CERN-wide AI platforms that support both technical systems and administrative functions

CERN will establish a common technical framework with safety and reliability criteria in Phase 1, progressing to cross-department AI platforms with centralised model registry, benchmarking, and continuous drift detection by Phase 3.

Key indicator: Standardised validation compliance across departments; increased reuse of AI components across systems; reduction in time-to-result.

Expected by: Phase 3.

Objective 2.3: Leverage AI to reduce operational costs, sustainability, and manual workload while maintaining or improving performance

CERN will pilot AI-assisted document processing, reporting support, and governed enterprise knowledge assistants in Phase 1, scaling to integrated workflow modules and CERN-wide knowledge management systems with structured reskilling programmes by Phase 4.

Key indicator: Measurable reduction in cycle time for selected processes; documented reduction in repetitive manual workload; maintained or improved user trust indicators.

Expected by: Phase 2 (pilots), Phase 4 (full deployment).

Objective 2.4: Apply strong governance, human oversight, and explainable AI to meet safety, security, and sustainability requirements

CERN will develop and operationalise a governance and risk framework for AI use in Phase 1, integrating AI models into standard system lifecycle management with monitoring, continuous improvement, and sustainability criteria by Phase 3.

Key indicator: Zero safety-critical or reputational incidents attributable to AI malfunction; documented and tested fail-safe mechanisms for all critical systems; increased regulatory and governance compliance.

Expected by: Phase 3 (framework operational), Phase 4 (sustained operations).

3.3 Pillar 3: AI as an Impact and Collaboration Enabler

CERN's role in the global scientific landscape extends well beyond its own experiments. As one of the world's leading scientific organisations, CERN has the opportunity, the responsibility, and the capability to contribute its AI knowledge to the broader scientific community, to AI-for-science initiatives in Europe and beyond, and to societal and industrial innovation through its Member and Associate Member States. This pillar addresses the external dimension of CERN's AI programme: how CERN positions itself as a reference site for AI-enabled science, how it transfers value to its stakeholders, and how it builds the human capital — inside and outside the Organisation — that the future of AI-driven research requires. CERN's commitment under this pillar is to translate internal AI capability into demonstrable external impact, guided by CERN's established Knowledge Transfer mandate and intellectual property policy for engagement with external actors beyond HEP. These external benefits arise primarily as a consequence of excellence in CERN's scientific and engineering mission rather than as an independent objective, and are pursued under clear governance so as not to divert focus or resources from the core programme, particularly through the demanding HL-LHC period.

Objective 3.1: Work closely with Member and Associate Member States, Europe, and industry to sustain CERN's leadership in AI-enabled science

CERN will map priority AI-for-science initiatives in EU, (A) MS and beyond and establish structured representation in selected flagship programmes in Phase 1, formalising long-term collaboration agreements and integrating CERN AI use cases into EU and international reference platforms by Phase 4.

Key indicator: Clear institutional visibility without mission drift; high-impact structuring partnerships in place, establish CERN convening role in AI development and deployment for science

Expected by: Phase 2 (engagement active), Phase 4 (formalised partnerships).

Objective 3.2: Translate AI advances from high-energy physics to societal, industrial, and innovation outcomes, ensuring return on public investments

CERN will define and consolidate its engagement strategy with the EU and (A)MS key AI-for-science initiatives in Phase 1, developing strategic partnerships and a clear value proposition for external collaborators, with periodic socio-economic impact assessments from Phase 3.

Key indicator: Documented pathways to socio-economic impact; clear institutional visibility without mission drift; high-impact structuring partnerships in place.

Expected by: Phase 3 (strategy consolidated), Phase 4 (impact demonstrated).

Objective 3.3: Integrate AI into education, training, and outreach to train and inspire future scientists, engineers, and administrators

CERN will build practical AI proficiency across its workforce, starting in Phase 1, covering the effective use of AI tools, chat interfaces, and agents and a sound understanding of their benefits and risks, developed through dedicated training courses, on-the-job learning, and awareness campaigns. More specialised and academic training will be delivered in collaboration with universities rather than through parallel in-house programmes, including doctoral and postdoctoral collaboration schemes, by Phase 4.

Key indicator: Increased AI competency coverage across CERN personnel; improved retention of AI-related talent; recognition of CERN as an attractive AI-for-science training environment.

Expected by: Phase 2 (framework active), Phase 4 (full coverage).

Objective 3.4: Contribute to Member and Associate Member States and European deep-tech and AI-for-science initiatives through knowledge transfer, industrial participation, and support of start-up creation

CERN will update its AI value proposition and establish focused industrial collaborative R&D in Phase 1, supporting start-up creation and scaling linked to CERN AI developments, with closer collaboration with national AI competence centres and periodic socio-economic impact assessments by Phase 4.

Key indicator: Measurable increase in AI-related knowledge transfer activities through existing channels; balanced industrial participation benefiting Member and Associate Member States; a proportionate contribution to the European deep-tech ecosystem within the programme's means.

Expected by: Phase 3 (initiatives launched), Phase 4 (impact sustained).

3.4 Lighthouse Initiatives

Having set out those objectives, the initiatives below show what they build towards. Each grows out of work already under way across the experiments, sectors, and departments, and is anchored to their detailed plans. They are presented in two tiers and ordered by maturity. The short-medium-term initiatives are considered to be achievable within the technology, time, and funding available to this plan at the conditions known today, and are the more mature group around which the plan is primarily organised. The horizon initiatives are deliberately more speculative and longer-term: they may begin during the plan but will mature mostly afterwards, in the Horizon 2030+, and depend on resources not available today, which CERN will pursue through its partnerships and collaborations programme. The set will be reviewed and adjusted as the programme progresses.

#	Lighthouse Initiative	What it delivers	Pillar / objectives	Delivery and horizon
Short-medium term initiatives — feasible within the plan's technology, time, and funding				
1	Common AI Infrastructure and Platform	Shared, CERN-wide AI compute, storage, and MLOps platform with a centralised model registry and full-lifecycle support, including agentic AI (MCP/LLM), serving scientific and operational AI across all sectors as common infrastructure for CERN use cases	Cross-organisational; CA2, Obj 2.2	Pilot Q1 2027; production 2028; operational across the plan
2	AskCERN: Information Discovery and Retrieval	LLM-based portal for discovery and retrieval across CERN-wide and domain-specific knowledge (built on curated corpora such as CDS, EDMS, or INSPIRE), improving productivity and streamlining processes and operations	Pillar 2; Obj 2.2, 2.3	Pilots from 2026; maturing through Phase 2; operational from 2028
3	Edge AI for Real-Time Data Acquisition	Shared capability for deploying validated AI models onto the hardware where the data is produced, covering FPGA and ASIC inference in trigger and readout chains and radiation-tolerant inference in accelerator equipment, with a common model-to-firmware workflow, benchmarking, and tooling reusable across experiments and accelerator systems	Pillar 1 / 2; Obj 1.1, 1.3, 2.1	Common tooling and benchmarks from Phase 1, building on work already under way in the experiments and in ATS; first shared reference deployments 2028; radiation-tolerant in-tunnel inference through Phase 4

4	Universal Fast Simulation Engines	Shared, community-validated AI-based fast and accurate simulation tools, building on AI fast-simulation tools already in production in some experiments, to replace the current fast-simulation software and complement and integrate with Geant4-based simulations, serving existing experimental collaborations and future-detector studies	Pillar 1; Obj 1.1, 1.3, 1.4	Building on tools already in production in some experiments, work begins in Phase 1; first results 2028; in production through Phase 4
5	AI-Driven Operations and Equipment Support	Agentic AI working alongside two communities: CCC operations teams (real-time HL-LHC tuning, automated setup/recovery, prioritised anomaly response) and equipment-group piquet/on-call teams (24/7 fault triage, natural-language retrieval over equipment docs and past interventions, drafting of intervention reports and logbook entries), plus a reusable governance template for safety-critical agentic AI (human-in/on-the-loop, hard guardrails, full auditability, staged autonomy) for reuse across the programme.	Pillar 1 / 2; Obj 1.2, 2.1, 2.3, 2.4	Operations agents 2029; piquet assistants piloted in Phase 2, scaled through Phase 3; governance template as of Phase 1
6	AI-Assisted Accelerator Systems Design	Generic, reusable platform that turns expensive first-principles simulations (ANSYS/MAPDL for thermal and mechanical-stress, CST for impedance, ngspice for circuits, Xsuite for beam dynamics, orchestrated at scale on HTCondor) into trained, uncertainty-aware surrogates, and wraps them in Bayesian/active-learning design optimisation and surrogate-based digital twins. Input into future related activities of detector design (LH-09).	Pillar 1; Obj 1.2, 1.3, 1.4	Simulation backends already operational in ATS; surrogate + optimisation layer piloted in Phase 2; surrogate-based digital twins operational by Phase 4
7	Physics Foundation Model	CERN-anchored foundation model built on CERN-hosted unique HEP datasets as shared scientific AI infrastructure, supporting tasks such as end-to-end reconstruction and simulation	Pillar 1; Obj 1.1, 1.3, 1.4	First models 2029; continued maturation into the Horizon 2030+
Horizon initiatives — speculative and long-term, pursued through partnerships and additional resources				
8	CERN AI Agent for Particle Physics	Agentic AI capable of assisting physicists in typical tasks related to experimental particle physics and phenomenology, such as performing a data analysis, assessing uncertainties, and learning the	Pillar 1; Obj 1.1, 1.4	Exploration from 2027 (Phase 2); substantial realisation in the Horizon 2030+, subject

		implications of physics results in terms of a given theoretical result		to additional resources
9	AI-Assisted Detector Design	End-to-end, differentiable optimisation of detector design, treating detector geometry, materials, and technology choices together with the reconstruction and analysis chain as a single objective that balances physics performance against cost, using differentiable programming and machine-learned surrogate simulators to explore the design space for next-generation detectors, in support of HL-LHC upgrades and future-facility studies such as the FCC, complementing the fast-simulation and physics-foundation-model initiatives	Pillar 1; Obj 1.1, 1.2	Exploration from 2027 (Phase 2); substantial realisation in the Horizon 2030+, subject to additional resources and collaborations
10	AI-Driven Material Discovery for Future Colliders	AI-driven discovery and inverse design of advanced materials for future-accelerator technologies, such as high-field superconductors and radiation-hard materials, using generative models and machine-learned property prediction to narrow the candidate space ahead of experimental validation, in support of FCC magnet and conductor R&D	Pillar 1; Obj 1.1, 1.2	Preparation from 2027–2028; substantial realisation in the Horizon 2030+, subject to additional resources and collaborations

Funds currently available at CERN allow it to establish the committed baseline described above and in the rest of this document. Building on this baseline, CERN aims to reach the more ambitious, horizon initiatives, such as the CERN AI Agent for Particle Physics, by securing additional funding and collaboration with AI research centres. These are stated as long-term ambitions, not commitments: they set a direction, motivate partnerships, and position CERN at the frontier, without drawing on the resources committed to the plan. The timeline below shows how each initiative is positioned across the phases and into the Horizon 2030+. The list of initiatives is indicative, not exhaustive and it might change as priorities, requirements, technology and resources evolve.

Lighthouse initiatives across the phased plan and the Horizon 2030+




Figure 3. Positioning of the lighthouse initiatives across the four phases of the plan and the Horizon 2030+. Short-medium term initiatives are delivered within the plan; horizon initiatives begin within it but mature mostly beyond 2030, subject to additional resources.

A detailed factsheet for each initiative, setting out its scope, expected impact, indicator, dependencies, resources, and milestones, is provided in the Lighthouse Initiative Factsheets section at the end of this document.

4. Phased Roadmap

The four-phase structure of this Implementation Plan is designed to sequence delivery in a way that manages risk, builds organisational capability progressively, and surfaces measurable outcomes at regular intervals rather than deferring value to the end of the programme. More details are given for the near term, Phase 1 is defined to quarterly granularity, while later phases are expressed at the level of outcomes and milestones. **The roadmap that follows is a snapshot at the time of writing: it reflects current requirements, technology, and expected outcomes, and will be monitored, reviewed, and adapted as these evolve. This makes the plan a living document, allowing flexibility for adjustment as the programme matures and lessons from early phases are incorporated.** The plan is monitored through periodic programme reviews by the AI Programme Committee and reviewed formally at least on an annual basis by the AI Programme Committee and CERN Management.

4.1 Phased Timeline

Phase 1 — Foundations (Q3 2026–Q1 2027)

Phase 1 establishes the organisational, technical, and governance foundations necessary to implement the AI Strategy at scale. Across four quarterly sub-phases, CERN activates its AI governance framework, deploys the first shared infrastructure and approved tools, puts risk and procurement controls in place, and launches priority pilots in scientific, operational, and

knowledge domains, and contributes, with the wider HEP community, to defining the major scientific challenges that can be addressed with AI in the HL-LHC and FCC time horizon. By the end of Phase 1 in Q1 2027, and subject to timely resourcing and recruitment, the organisation aims to have a functioning governance model, a baseline AI compute and development environment, a Responsible AI Policy, a portfolio management process, and active AI communities of practice, positioning all sectors to begin scaling from Q2 2027. It has also a vision for future breakthrough potential and a mechanism to flexibly adapt to the rapid pace of AI development and revise this plan accordingly, capitalising on the collective bottom-up ideation and innovation culture that constitutes one of the major strengths of CERN and the HEP community.

Q2 2026: The Implementation and Resource Plan is formally approved by CERN Management, establishing the authorised framework within which all subsequent actions are planned and resourced.

Q3 2026: AI governance becomes operational with the establishment of the AI Programme Committee and its Terms of Reference; the first CERN-approved AI tools are provisioned organisation-wide; priority scientific AI workstreams are reassessed and continued across RCS; and CERN's external positioning on AI is activated through EC, EIROforum, and international engagement.

Q4 2026: A pilot compute pool becomes operational; the ATS Core Team is assembled and pilot workstreams for predictive maintenance and AI assistants are launched; AI communities of practice are established across the organisation; portfolio management for coordination and shared awareness, building on existing reporting, is integrated into annual planning; the MLOps pilot platform and AI development environment baselines are delivered; and an ATS-wide workshop consolidates and aligns accelerator AI activities ahead of 2027 scaling, while an ATS-RCS domain workshop on common areas defines co-development opportunities.

Q1 2027: The ethical risk classification and procurement principles frameworks are in place; the AI Infrastructure Strategy is approved; and the CERN Responsible AI Policy, developed through consultation with the community, is approved.

Phase 2 — Scale-up and Integration (Q2–Q4 2027)

Phase 2 transitions the foundations established in Phase 1 into operational capabilities. Data governance frameworks, sector data maturity assessments, and standardised data connectors are put in place; the hybrid AI platform reaches pilot production; and the Responsible AI Policy is fully integrated into CERN governance. Competency frameworks, advanced training tracks, and the AI monitoring and assessment framework, covering resource deployment and programme impact and consolidated from sector reporting, become operational, enabling the organisation to manage production AI systems with confidence. By the end of Phase 2, CERN moves from a governance and infrastructure posture to one of active, managed AI deployment across multiple domains. On the scientific side, the first lighthouse initiatives move from pilot towards early results, with work on the Universal Fast Simulation Engine scaling alongside the data and platform foundations. In parallel, exploratory work begins on the horizon initiative CERN AI Agent

for Particle Physics, pursued through partnerships and additional resources rather than the committed programme.

Phase 3 — Production and First Outcomes (2028)

Phase 3 marks the transition from capability-building to demonstrable results. The MLOps platform reaches full production and model documentation standards, proportionate to the type and risk of each system, are applied across production AI systems. The first annual AI compliance assessment, conducted by end 2027, provides the organisation with its first structured view of AI governance maturity. The assessment is integrated into existing institutional review cycles to avoid additional reporting overhead. Data stewardship roles are operational, and priority datasets are catalogued. First results from the programme's AI workstreams are delivered by RCS in pattern recognition, fast simulation, and online processing, building on CERN's established use of machine learning in these areas, and ATS pilot systems transition towards operational deployment. By the end of 2028, CERN has moved AI from a programme under construction to a programme delivering measurable value. Scientifically, the first lighthouse outcomes are expected in this phase, with the Universal Fast Simulation Engine delivering community-validated results that feed detector design and analysis.

Phase 4 — Full Production, HL-LHC Readiness and Sustained Operations (2029–2030)

Phase 4 delivers the full operational maturity of CERN's AI capability across all domains and ensures alignment with the HL-LHC programme and the longer-term FCC preparation. The shared AI platform reaches full production with continuous validation and benchmarking; agentic AI assistants and production-grade ML infrastructure are deployed at scale; scientific AI tools already used in the experiments are more deeply integrated and scaled into HL-LHC workflows, and new tools are designed with FCC requirements in mind; and predictive maintenance, real-time optimisation, and quality assurance systems operate across accelerator facilities. Workforce AI literacy achieves full coverage. Continued oversight by the AI Programme Committee and annual review cycles, building on the reviews already performed within the sectors, maintain alignment with CERN's mission and strategic objectives through 2030 and beyond. The major scientific initiatives mature here: the CERN Physics Foundation Model and AI-Driven Operations and Equipment Support (2029), and the CERN AI Agent for Particle Physics (2029 to 2030), all aligned with the HL-LHC programme and the FCC preparation. In this context, HL-LHC readiness means these activities are fully under way by 2030.

Beyond the Plan, Horizon 2030+

This outlook lies outside the scope and budget of the present plan. It is included to show how the concrete capabilities delivered within the five years, the common AI infrastructure, the foundation models, and the validated scientific and operational tools, mature towards the FCC era. With additional resources secured through partnerships and collaborations, these foundations support the more speculative, horizon initiatives, principally the CERN AI Agent for Particle Physics and the continued maturation of the physics foundation models. Nothing in this

phase is a commitment of the current programme: it depends on funding and partnerships still to be established, and it frames the case CERN will make for them.

5. Governance

The governance of this Implementation Plan operates under the principle of *freedom in a framework* adopted in the CERN AI Strategy and established by the CERN CIO governance model. Responsibility for scientific direction, operational execution, and local implementation of AI solutions remains with sectors, departments, experiments, and services. CERN-wide governance defines shared principles, priorities, and decision mechanisms that enable local autonomy while avoiding fragmentation, unmanaged risk, and duplication of effort. Governance arrangements, including the roles of the C4 and the AI Programme Committee, are reviewed transparently on an annual basis, with consultation of the sectors and the scientific community, and updated as the programme evolves.

Roles and responsibilities. CERN Management holds ultimate accountability for the Implementation Plan. It approves strategic priorities and the resource envelope, endorses AI as a core organisational capability, and reviews progress, risks, and value creation on an annual basis. The **CERN CIO** chairs the CIO Committee for Coordination and Collaboration (C4), which steers and coordinates the implementation of the AI Strategy across the Organisation as the main steering body for digital transformation initiatives, and supports cross-organisational coherence, coordinates and arbitrates resourcing of the plan based on recommendations from the AI Programme Committee, and is the primary point of accountability for reporting to CERN Management. A senior leader within the CIO Office with responsibility for AI Strategy oversees the programme and the mobilisation of resources, maintains the AI portfolio dashboard, monitors progress against milestones, and chairs the AI Programme Committee. The **AI Programme Committee (AIPC)** is the principal cross-organisational body for coordination, monitoring, and recommendation: it **evaluates** opportunities where CERN can provide innovative development of AI technologies and capabilities, supports the prioritisation of AI initiatives and the adoption of shared standards and platforms, and makes recommendations to the C4 for decisions and arbitration requiring management approval. **Sectors and departments** retain full responsibility for their share of the cross-organisational and sector-specific actions, and report progress in activities, technical results, impact, and financial expenditures to the AI PC on a regular basis.

Decision escalation. Decisions are taken at the lowest appropriate level consistent with the scope and risk of the action. Technical and operational decisions are taken by the sector teams. Decisions involving significant resource reallocation or cross-sector dependencies are escalated to the C4. Decisions requiring CERN Management endorsement — including the annual portfolio and resource review — are submitted by the CIO following C4 deliberation.

Responsible AI and risk management. All initiatives are subject to the risk classification model and Responsible AI framework integrated in the CERN Risk Register maintained by OSI. Safety-critical and mission-critical applications require mandatory formal review before deployment, applied to CERN-operated and institutional AI systems, while scientific applications governed by the experiment collaborations remain under their own review processes, with defined fail-safe mechanisms and escalation paths. An incident reporting mechanism and annual AI compliance

assessment provide ongoing assurance, with the first assessment due by end 2027. Exploratory AI development in pilot contexts is actively encouraged, provided appropriate safeguards are in place and lessons are systematically shared. The detailed risk assessment, scoring methodology, and risk register is maintained by the AI PC and contributed to the global CERN Risk Register.

Portfolio management and monitoring. The AI portfolio is managed through a centralised dashboard maintained by the AI PC and reviewed at least annually by the C4 and CERN Management. Prioritisation follows the principles established in the AI Strategy; initiatives that cannot demonstrate alignment with strategic objectives may be deferred, rescope, or discontinued. The AI PC reports quarterly to the C4 through its Chair, and prepares an annual report covering delivery progress, resource utilisation reported by the sectors and departments and consolidated by the AIPC, risks, and any proposed adjustments. The Implementation Plan itself is subject to an annual review, with material revisions submitted to CERN Management for approval.

The CERN AI programme governance model and composition are described on the AI Programme Committee section of the CERN CIO Artificial Intelligence web site³.

6. Resource Plan

The Resource Plan translates the implementation actions defined in Section 6 into a financial and human effort envelope for the period 2026–2030 (with some effects on the 2031 calendar year). It covers costs in two categories — personnel across all profiles, and infrastructure including computing, software, services, and licensing — and is structured to distinguish clearly between three funding sources: existing investments already committed within each sector or department; external funding through EC projects, foundations, or industrial collaborations aligned with CERN's AI objectives; and the additional CERN investment required to scale from current pilot activities to sustained production capability. The plan is also structured over time through annual estimates, allowing resource commitments to be reviewed and adjusted in line with delivery progress and evolving priorities.

It should be noted that externally funded projects carry committed objectives and timelines that cannot be unilaterally changed. These projects must be brought to completion according to their existing terms, aligning activity to the AI Strategy wherever possible within those constraints. This applies to externally funded projects such as NGT and EU-funded projects, which carry their own milestones and deliverables, as well as to activities embedded within scientific collaborations and their governance structures, such as the EP R&D programme and the LHC experiments. The same consideration applies to existing CERN-funded resources, which remain under departmental management and may be partially redirected to higher-priority tasks. The resource figures in this plan are a snapshot of current and projected resources rather than a commitment, and should be read as such.

³ <https://cio.web.cern.ch/artificial-intelligence/>

7. Implementation and Operations

The strategic objectives set out in Section 3 are delivered through two complementary categories of action, both defined in detail in the detailed workplans managed by the AIPC and the CERN Sectors/Departments.

Cross-organisational structuring actions (CA1–CA5) establish the shared foundations on which all sector-level implementation depends. They address governance and decision frameworks, infrastructure and platforms, data foundations, workforce and competency, and responsible AI. These actions are CERN-wide in scope, coordinated by the CIO Office and the AI Programme Committee, with execution remaining within the sectors and departments, and are sequenced to be operational before or in parallel with the sector-specific work they enable.

Sector-specific implementation plans translate the pillar objectives into prioritised actions, deliverables, and indicators for each of CERN's five sectors: Accelerators and Technology (ATS), Research and Computing (RCS), Site Operations (SO), Finance and Human Resources (FHR), and Stakeholder Relations (SRS). Each sector retains full execution responsibility for its plan, operating within the shared framework established by the cross-organisational actions and reporting progress to the AI Programme Committee on a regular basis. The sector plans are set out in full in the yearly Workplans of each department.

Additional actions may emerge through opportunities for cross-sector collaboration or co-development with the HEP community, for example in algorithm development, robotics, or shared software tools. Such opportunities will be supported on an ongoing basis through the AI Programme Committee's portfolio management process.

Lighthouse Initiative Factsheets

Common AI Infrastructure and Platform

Initiative ID	LH-01	Tier	Short-medium term
AIPC Lead	AI Infrastructure Lead	Sector / dept	RCS-IT
Pillar	Cross-organisational	Objectives	CA2, Obj 2.2

1. WHAT IT DELIVERS

Shared, CERN-wide AI compute, storage, and MLOps platform with a centralised model registry and full-lifecycle support, including a training platform, agentic AI (MCP/LLM), and assistants, serving scientific and operational AI across all sectors as common infrastructure, including in the Technical Network.

2. SCOPE

In scope

Platform with advanced functionality for LLMs and Agentic AI (MCP servers, tools, prompts, skills), covering discoverability, access control, orchestration, observability.
Support for both the general (GPN) and technical network (TN).
Support for both on-premises and external/cloud hosted models.
Provisioning of coding assistants and AI chatbots.
Long term strategy definition for access to accelerator resources, in particular GPUs, and including hybrid deployments (HPC and cloud) and evolution of datacentre infrastructure.
Documentation and dissemination material with best practices for building tools and applications around the supported toolchain.

Out of scope

Development of MCP servers for the different CERN services - expecting these to come built-in with off-the-shelf tools (GitLab, Atlassian, etc) and from each team for CERN specific tools (LanDB, CDS, Indico, etc).
Ownership and curation of knowledge bases.
Coverage of license costs for coding assistants and AI chatbots.

3. EXPECTED IMPACT

Consolidation of hardware resources in a centralized pool for optimized allocation and utilization - especially relevant given the scarcity, high cost, power and cooling requirements of accelerator hardware, in particular GPUs.

Reduction of overall cost of supporting use cases (hardware and person power), both at the service level and application teams.

Long term sustainability of the platform, enabling adequate support and service maintenance, consolidating the stack used by our communities and aligning with industry standards.

A reference architecture for a full featured, scalable platform supporting AI and ML for scientific computing. This is expected to be reusable in large extent by other sites supporting HEP workloads (WLCG) as well as industry organizations targeting similar use cases. The [first version](#) has already been published in the cloud native computing foundation (CNCf).

Indicator: Overall allocation and utilization metrics of hardware resources. Performance metrics (requests and tokens per second and CHF) achieved for key use cases vs expected/published by vendors and other organizations.

4. MILESTONES

Quarter	Phase	Description	Type
2026-Q4	Phase 1	Platform architecture and service definition agreed, covering tenancy, MLOps toolchain, model registry and support model.	MTP
2027-Q1	Phase 1	Platform availability for initial use cases, covering optimized LLM inference as well as early support for MCP servers, tools and prompts, agents and skills. Includes discoverability, orchestration and observability of all components.	MTP
2027-Q1	Phase 1	Procurement completed for the initial hardware budget allocation, covering the bulk of resource needs across all initiatives. Target hardware delivery in Q3 2027, with significant uncertainty regarding this timeline due to the current high demand for accelerator resources (GPUs).	MTP
2027-Q3	Phase 2	Dissemination and training material for communities integrating tools and applications with the platform. Regular schedule for training events and office hours. The timeline for this milestone is highly dependent on the final hardware delivery timeline.	MTP
2027-Q4	Phase 2	Generally available and production-ready service, including resource lifecycle management and day-2 validation of business continuity and disaster recovery. Support for hybrid deployments (cloud and HPC). The timeline for this milestone is highly dependent on the final hardware delivery timeline.	MTP
2028-Q4	Phase 3	Platform in production, serving scientific and operational AI across sectors.	MTP
2030-Q2	Phase 4	Continuous validation and benchmarking in operation, capacity aligned to HL-LHC needs.	MTP

5. DEPENDENCIES

Enabling platform for LH-02, LH-03, LH-05, LH-06, LH-07. Depends on CA2 and on IT department capacity planning.

6. RESOURCES, RISKS AND STATUS

Resources, existing	Resources, new	Top risks, max 3
NGT 1.1 person power supporting the MLOps platform and toolchain (until 2027, partially in 2028). Very limited available and suitable hardware resources in IT (H100 or newer) - existing resources largely allocated to other use cases.	Person power beyond NGT from MTP budget, covering the platform and service activities for both traditional ML and LLM/Agentic AI use cases. Bootstrap hardware procurement in IT prior to MTP allocation (2026). Large scale procurement from MTP budget (2026/27).	Continuation of cost volatility and long delivery timelines of suitable accelerator hardware, in particular GPUs. Trends in accelerator hardware (power and cooling needs) vs capabilities of our existing on-premises infrastructure. Requirement of liquid cooling for post-Blackwell generations of Nvidia GPUs. Stack for LLM and Agentic AI remains a fast-moving target - MCP, Skills, and many new developments on the inference space.

Status: Not yet rated **Last update:** 31 August 2026 **Next review:** AIPC plenary, September 2026

AskCERN: Information Discovery and Retrieval

Initiative ID	LH-02	Tier	Short-medium term
AIPC Lead	AI for Productivity Lead	Sector / dept	All CERN Sectors
Pillar	Pillar 2	Objectives	Obj 2.2, 2.3

1. WHAT IT DELIVERS

LLM-based portal for knowledge discovery and retrieval across CERN-wide and domain-specific knowledge (built on curated corpora such as CDS, EDMS, or INSPIRE), improving productivity and streamlining processes and operations.

2. SCOPE

In scope

Proof of Concept (PoC) testing 3–4 iconic use cases in a knowledge-retrieval pilot aligned with CERN's strategy pillars.

User evaluation of the PoC, based on defined success criteria and evaluation metrics (incl. cost vs business value, technical performance, and user adoption). Comparison against alternative solutions for management decision-making.

Design, implementation and deployment of a production AskCERN solution, that can be scaled up and applied to multiple differing internal and external facing CERN knowledge sources.

Cataloguing of secondary use cases for AIPC prioritization. Supporting use-case owners CERN-wide to self-serve production deployment on the production platform.

Out of scope

General-purpose web search AI assistant capabilities beyond CERN content/knowledge.

Development of MCP servers for the different CERN services - expecting these to come built-in with off-the-shelf tools (GitLab, Atlassian, etc) and from each team for CERN specific tools (LandB, CDS, Indico, etc).

3. EXPECTED IMPACT

Improve access to organisational knowledge, while also reducing effort spent searching for information and answering repetitive queries.

Demonstrate value of generative AI in strategic CERN business domains, and provide evidence to inform future AI investment decisions.

Establish reusable architecture to leverage AskCERN-type AI solutions beyond the initial 3-4 use cases tested in the PoC.

NB: there is limited scope for external impact and exploitation. AskCERN may answer the needs of external partners with CERN (e.g. finding info on doing business with CERN, candidates applying to CERN asking questions about CERN vacancies).

Indicator: number of evaluated answers are judged accurate and grounded (answer quality); number of users rate answers as useful (user satisfaction); reduction in time spent finding information (time saved); reduction in repetitive information requests to support teams (reduction in support effort); number of users use AskCERN at least weekly/daily (adoption), cost in tokens to achieve a specific task (cost/value).

4. MILESTONES

Quarter	Phase	Description	Type
2026-Q4	Phase 1	Testing 3-4 use cases in a knowledge-retrieval pilot targeting a well-defined user community on curated corpora, measuring value to CERN and scalability.	MTP
2027-Q2	Phase 2	Corpus coverage extended across CERN-wide and domain-specific sources, retrieval quality baseline established.	MTP
2027-Q4	Phase 2	Access control and data protection review completed, service hardened.	MTP
2028-Q2	Phase 3	Operational service with a defined support model.	MTP
2029-Q4	Phase 4	Adoption and productivity indicators reported to the AIPC.	MTP

5. DEPENDENCIES

Corpus owners (CDS, EDMS, INSPIRE, Admin eGuide, HR website(s), other admin procedure websites e.g. Procurement, Safety, etc.). Data protection and Legal review. Serving capacity from LH-01. PoC use case aligned with LH-05.

6. RESOURCES, RISKS AND STATUS

Resources, existing	Resources, new	Top risks, max 3
Project team capacity Stakeholder capacity for making knowledge corpora available in Sectors	Ca. 200-300k CHF for the 2026 PoC Budget for ongoing AskCERN service from Jan 2027, TBC. Estimated 850k EUR to over 2.5m EUR for 3 years (Price Enquiry required), not part of the MTP request	Costing models of the AI tested may be beyond CERN's budget allocated for AskCERN. Upholding CERN Procurement rules, Data Privacy or security constraints may delay the purchase of the service for the PoC and the eventual implementation of a production AskCERN service. Stakeholders may not be available for making knowledge corpora available, for workshops or testing.

Status: Not yet rated **Last update:** 31 August 2026 **Next review:** AIPC plenary, September 2026

Edge AI for Real-Time Data Acquisition

Initiative ID	LH-03	Tier	Short-medium term
AIPC Lead	AI for Research Lead	Sector / dept	RCS
Pillar	Pillar 1 / 2	Objectives	Obj 1.1, 1.3, 2.1

1. WHAT IT DELIVERS

Shared capability for deploying validated AI models onto the hardware where the data is produced, covering FPGA and ASIC inference in trigger and readout chains and radiation-tolerant inference in accelerator equipment, with a common model-to-firmware workflow, benchmarking and tooling reusable across experiments and accelerator systems.

2. SCOPE

In scope Development and maintenance of an Edge-AI software ecosystem for hardware-aware training and on-edge deployment, meeting current and future needs for LHC operations and CERN detectors. Support for emerging hardware technologies beyond FPGA (AI engines, etc.). Development of a test platform for inference accuracy in a radiation environment. Meet requirements towards adoption by the experiments. Support for software-to-hardware pipelines enhancing technological sovereignty of CERN member states. Share knowledge on edge-AI deployment with stakeholders, through participations in research initiatives beyond HEP, with emphasis on socially impactful application (ocean plastic detection, satellite debris, fire detection through geo observation)	Out of scope Share knowledge and promote the use CERN edge-AI solutions for applications external to CERN mission (this requires additional support for knowledge transfer activities from the cross-organisational initiatives) Adoption by experiments
---	--

3. EXPECTED IMPACT

CERN experiments and accelerator complex are already relying on on-edge AI solutions. By promoting a research program on edge AI, this lighthouse project aims at sustaining this effort and its development towards new technologies and use cases, e.g., the design of AI ASICS and the deployment on front-end electronics (both for accelerators and experiments).

Indicator: Number of supported model-to-firmware pipeline solutions and covered backend use cases

4. MILESTONES

Quarter	Phase	Description	Type
2026-Q4	Phase 1	Assess state of the art; identify CERN possible contribution, establish link to existing HEP community efforts (e.g., FastML collaboration) to discuss developments in the context of ongoing activities	MTP
2027-Q1	Phase 1	Assessment of supported model-to-firmware pipeline solutions and covered backend use cases. Identification of existing and missing ingredients for	MTP

Quarter	Phase	Description	Type
		frontend use-case support (ASICs and rad-hard FPGAs). Exploration of synergies with external partners (e.g., ESA for satellite applications).	
2027-Q4	Phase 2	Roadmap of post-NGT development for backend and frontend solution support, including support for emerging hardware solutions, synergy with EP R&D program and regional initiatives.	MTP
2028-Q4	Phase 3	Assessment of NGT legacy and post-NGT work planning, including synergies between detector and accelerator use cases.	MTP
2029-Q4	Phase 4	Demonstration of radiation-tolerant inference, suitable for deployment in the LHC tunnel and detector caverns.	MTP
2030-Q4	Phase 4	Assessment of HL-LHC accelerator and detector adoption for Run 4 startup. Planning future R&D program, including support for non-LHC experiments.	MTP

5. DEPENDENCIES

Experiment L1 trigger upgrade schedules for HL-LHC. Next Generation Triggers, to end 2028. ATS in-tunnel deployment programme. Model management and training capacity from LH-01.

6. RESOURCES, RISKS AND STATUS

Resources, existing	Resources, new	Top risks, max 3
NGT resources until 2028	Continuation of NGT activities in this area, partially in the MPT request. Possibility to discuss continuation of NGT.	No NGT 2 or equivalent funding might materialise, funds not enough

Status: Not yet rated **Last update:** 31 August 2026 **Next review:** AIPC plenary, September 2026

Universal Fast Simulation Engines

Initiative ID	LH-04	Tier	Short-medium term
AIPC Lead	AI for Research Lead	Sector / dept	RCS
Pillar	Pillar 1	Objectives	Obj 1.1, 1.3, 1.4

1. WHAT IT DELIVERS

Shared, community-validated AI-based fast and accurate simulation tools, building on AI fast-simulation tools already in production in some experiments, to replace the current fast-simulation software and complement Geant4-based simulations, serving existing experimental collaborations and future-detector studies.

2. SCOPE

In scope	Out of scope
Develop AI-driven fast simulation training and deployment solutions that can generalize across experiments. Develop universal architectures for AI-driven detector simulations (e.g., through foundation models) and inference engine to speed up GEANT4 simulation. Creation of a benchmark platform, allowing for a fully comprehensive performance comparison for input models. Support the acceleration of event generators through the development of AI-based solutions and surrogate models. Meet requirements towards adoption by the experiments Develop broad-scope software solutions for simulation tuning and deployment, to also meet requirements from non-HEP GEANT4 users	Adoption by experiments

3. EXPECTED IMPACT

Event generation and simulation is a major consumer of computing resources in HEP. By accelerating these workflows with the AI solutions, this lighthouse aims at providing a workflow which could deliver equal accuracy with less time and less computing resources. Provided its adoption is achieved, this lighthouse could support meeting the computing requirements for HL-LHC and future experiments.

Indicator: To be defined after completion of initial assessment

4. MILESTONES

Quarter	Phase	Description	Type
2026-Q4	Phase 1	Assess state of the art; identify CERN possible contribution, establish a HEP community forum (horizontal across theory groups and experiments) to discuss developments in the context of ongoing activities	MTP
2027-Q2	Phase 1	Establishment of an R&D roadmap towards a universal architecture and training pipeline, including both fast and flash simulation solutions	MTP
2027-Q4	Phase 2	Definition of benchmark datasets	MTP
2028-Q4	Phase 3	Reference implementation available, trained on the scope of existing production tools.	MTP

Quarter	Phase	Description	Type
2029-Q4	Phase 4	First community-validated results delivered, feeding detector design and analysis.	MTP
2030-Q4	Phase 4	In production use for HL-LHC workflows and future-detector studies.	MTP

5. DEPENDENCIES

Adoption depends on the experiment collaborations. Training capacity from LH-01.

6. RESOURCES, RISKS AND STATUS

Resources, existing	Resources, new	Top risks, max 3
Ongoing work in EP-SFT, supported until 2028 by the EP R&D programme, and contributions from the community	Effort supported by the MTP allocation, contributions from the R&D WG on GEN-AI solutions for FastSim	Universal architecture might not generalise as expected; Experiments may not have (or invest) resources to adopt the proposed solutions

Status: Not yet rated **Last update:** 31 August 2026 **Next review:** AIPC plenary, September 2026

AI-Driven Operations and Equipment Support

Initiative ID	LH-05	Tier	Short-medium term
AIPC Lead	AI for Accelerators Lead	Sector / dept	ATS / BE-SY-EN-TE, RCS
Pillar	Pillar 1 / 2	Objectives	Obj 1.2, 2.1, 2.3, 2.4

1. WHAT IT DELIVERS

Operational AI for the ATS complex, delivered by the core team developed alongside the groups leveraging the CERN AI implementation plan resources:

- On the equipment side, automated anomaly detection, early fault prediction and prescriptive maintenance that catch degradation before it causes downtime, building on the prototypes already running in the SY groups (ABT, RF, BI, EPC) and related work in TE (CRG, VAC, MPE) and extending them to further equipment and to the main infrastructure groups.
- On the operations side, first-line assistants for equipment piquet, operations and software support, and agents working alongside the CCC and experts (real-time machine tuning, automated setup and recovery, prioritized anomaly response, drafting of intervention reports and logbook entries...).

These solutions are delivered on a small set of shared enablers, so the work is not repeated group by group: a common approach to detection and prescriptive maintenance, an agentic standardization on the CERN IT LLM service (LH-01), a shared evaluation system for comparing models and architectures, and a reusable governance template for safety-critical agentic AI.

The core team builds and maintains the shared enablers; the groups own their device integration, models, playbooks and the day-to-day operation of the deployed solutions.

2. SCOPE

In scope

The shared substrate the solutions run on and its standardization (the agentic layer on the CERN LLM service, the shared evaluation system and the safety-critical governance template), together with hands-on support to the groups up to deployment.

The data, domain models, integration, playbooks and operations that turn that substrate into working equipment-health and operations-support solutions.

Coverage extends the existing equipment work (ABT, RF, BI, EPC, CRG, VAC, MPE and EPA) to further equipment and infrastructure groups, and to the piquet, operations, software-support and CCC functions.

AI stays advisory or supervised, human in or on the loop, and the enabling framework for each use case is settled through a technical review of the options.

Out of scope

Unsupervised autonomy: no agent acts outside an approved scope or without a defined human-in- or on-the-loop mode. No changes to machine protection or interlock systems. The LLM service itself is provided by IT, and the platform, data and MLOps substrate come from CA1–CA5. The federated-learning platform adaptation for CERN is adopted from there where a use case needs it; surrogate modelling and design optimization sit in LH-07. The groups keep ownership of their device integration, models, playbooks, thresholds and the operation of the deployed solutions.

3. EXPECTED IMPACT

Efficient use of Expert time, in the CCC and on equipment piquet, the scarcest resource around the complex. One maintained framework per pattern, rather than a stack per group, should cut unplanned downtime by catching degradation before failure, shorten fault diagnosis, absorb

routine and 24/7 load, and make first-line response consistent regardless of who is on shift. Directional: targets to be agreed with the groups as the frameworks reach production.

Indicator: Number of equipment and infrastructure groups running on the shared frameworks and number of systems in production; reduction in unplanned downtime and in first-line fault resolution time against a pre-framework baseline; agent task-success rate on an ATS benchmark; number of safety or access-control incidents.

4. MILESTONES

Quarter	Phase	Description	Type
2027-Q1	Phase 1	Governance template for safety-critical agentic AI issued, covering human-in and on-the-loop, guardrails, auditability and staged autonomy.	MTP
2027-Q2	Phase 1	Technical assessment of the available frameworks completed, identifying gaps and setting the roadmap for the substrate to build on.	MTP
2027-Q3	Phase 1	Agentic framework, first version running on the CERN LLM service, with governed connectors, skills over the main CERN knowledge sources and a shared evaluation system.	MTP
2027-Q4	Phase 1	Prescriptive maintenance framework, beta version, available to identified groups (equipment groups and the main infrastructure groups).	MTP
2027-Q4	Phase 2	Piquet assistant piloted with at least one equipment group.	MTP
2028-Q4	Phase 2	Both frameworks at first stable version: prescriptive maintenance in production across the participating equipment and infrastructure groups, and first-line assistants for piquet scaled to further equipment groups, operations and software support in supervised use.	MTP
2029-Q2	Phase 3	Both frameworks tested on the first machines and equipment returning to operation after LS3, with the results feeding back into the framework releases.	MTP
2030-Q4	Phase 4	Operations agents in supervised use in the CCC for real-time tuning and automated setup and recovery; staged autonomy reviewed with operations.	MTP

5. DEPENDENCIES

The CERN LLM service (LH-01), with API access to state-of-the-art models, is a mandatory dependency for the agentic framework.

Prescriptive maintenance requires a unified and standardized access layer to Controls equipment and supervisory systems, aligned with the Next Generation Controls Program (specifically NGC WP2).

6. RESOURCES, RISKS AND STATUS

Resources, existing	Resources, new	Top risks, max 3
Ongoing pilots and tests	Effort partially funded by the MTP, needs additional funds from external projects	If no additional funds become available, the main risk is that the solutions are not operationalised

Status: Not yet rated Last update: 31 August 2026 Next review: AIPC plenary, September 2026

AI-Assisted Accelerator Systems Design

Initiative ID	LH-06	Tier	Short-medium term
AIPC Lead	AI for Accelerators Lead	Sector / dept	ATS
Pillar	Pillar 1	Objectives	Obj 1.2, 1.3, 1.4

1. WHAT IT DELIVERS

AI models of accelerator hardware and beams that make design faster and give operations a live view of systems that cannot be measured directly; leveraging a common framework built and maintained by the ATS AI core team on the groups' requirements and delivered with the resources the AIP and groups commit.

The framework lets an expert wrap a model, be it a simulator, a neural network or a reduced one, as a virtual device that runs alongside the real machine, calibrated against it and deployed in operation.

- Surrogate-based design optimization: fast, uncertainty-aware approximations of expensive first-principles simulations driven by Bayesian optimization and active learning, turning loops that take days into interactive ones.
- Generalization of the emulation and hardware-optimization work already started in ATS covering geometry with point-cloud networks, mechanical and thermal properties, impedance and beam characteristics, on top of the high-fidelity backends already orchestrated at scale on HTCondor.

Where the models are validated, differentiable simulation over the beam-dynamics codes opens gradient-based optimization, calibration and sensitivity analysis. Virtual sensors infer quantities the machine does not measure (for example 3D temperature in accelerator equipment), and digital twins of installed systems (i.e. LHC superconducting circuits with ageing models) give what-if and predictive studies next to the running machine.

2. SCOPE

In scope

Solutions shared enablers: a surrogate-based design-optimization platform, and a framework that:

- lets an expert wrap a model (simulator/ a neural network or a reduced one) as a virtual device that runs alongside the real machine, calibrated against it and deployed in operation.
- offers uncertainty quantification, out-of-distribution detection and recalibration, for an increased surrogates robustness.

Out of scope

The high-fidelity simulation codes themselves: ANSYS, CST, ngspice and Xsuite stay with their owners, as do the licences. No replacement of first-principles simulation where accuracy is the requirement; surrogates are used inside their validated range. Compute and MLOps come from CA1–CA5. Models and design cases stay with the groups. Prescriptive maintenance and agentic assistants sit in LH-05.

3. EXPECTED IMPACT

Faster, better-informed design and clearer sight of installed systems. Simulation loops that take days become interactive, so engineers explore far more of the design space and freeze designs later and on better evidence, with fewer expensive high-fidelity runs; the effect is largest for FCC equipment still on the drawing board. In operation, virtual sensors expose quantities that are not directly measured, and ageing-aware digital twins support predictive and what-if studies next to the running machine, informing decisions without touching it. Built once and reused across groups, the platform also cuts duplicated surrogate work across the sector. Targets are directional until each model class is validated with the owning group.

Indicator: Number of model classes and groups with validated surrogates on the platform; accuracy against first-principle results within the stated uncertainty; reduction in high-fidelity runs and in wall-clock time per design iteration; number of design studies and of virtual devices in operation.

4. MILESTONES

Quarter	Phase	Description	Type
2027-Q2	Phase 2	Platform at first release: containerised simulation backends for circuits, thermal-mechanical and beam dynamics, driven by declarative study definitions and orchestrated on HTCondor.	MTP
2027-Q4	Phase 2	Surrogate and optimisation layer on one model class: training, uncertainty quantification and Bayesian optimisation, validated against first-principle results for a reference case.	MTP
2028-Q4	Phase 3	Platform extended to geometry with point-cloud networks, mechanical, thermal, impedance and beam dynamics, with gradient-based optimisation where simulations are differentiable; digital-twin framework at first version.	MTP
2029-Q4	Phase 4	Series of virtual devices deployed for one accelerator zone; digital-twin framework at first stable version. Platform used end to end in at least two design studies in different groups.	Ext. Funds
2030-Q4	Phase 4	Virtual devices in operation next to the machine, calibrated against live conditions.	Ext. Funds

5. DEPENDENCIES

Continued availability of the simulation backends and of the licences for the commercial codes. Compute orchestration at scale (HTCondor). Controls integration to run virtual devices alongside the machine. Group ownership of the models and design cases.

6. RESOURCES, RISKS AND STATUS

Resources, existing	Resources, new	Top risks, max 3
Ongoing pilots and tests	Effort partially funded by the MTP, needs additional external resources	If no additional funding becomes available, the main risks is that the solutions cannot be operationalised

Status: Not yet rated **Last update:** 31 August 2026 **Next review:** AIPC plenary, September 2026

Physics Foundation Model

Initiative ID	LH-07	Tier	Short-medium term
AIPC Lead	AI for Research Lead	Sector / dept	RCS
Pillar	Pillar 1	Objectives	Obj 1.1, 1.3, 1.4

1. WHAT IT DELIVERS

CERN-anchored foundation model built on CERN-hosted unique HEP datasets as shared scientific AI infrastructure, supporting tasks such as end-to-end reconstruction and simulation.

2. SCOPE

In scope	Out of scope
Creation of the largest scientific dataset in the world, consisting of HL-LHC collisions simulated in various detectors. Definition of a generalizable foundation model to process LHC collisions, including broad-scope tokenization strategy. Creation of a benchmark platform, allowing for a fully comprehensive performance comparison for input models. Meet requirements towards adoption by the experiments	Adoption by experiments

3. EXPECTED IMPACT

Foundation models are multi-purpose models that can be tuned to multiple tasks. The groundbreaking impact of these models is testified by the impact that LLM agents had on many areas, including science. This lighthouse aims at developing a large, accurate and generalizable model that could be used as a baseline to develop AI-based global event descriptions.

Indicator: To be defined after initial assessment

4. MILESTONES

Quarter	Phase	Description	Type
2026-Q4	Phase 1	Assess state of the art; identify CERN possible contribution; and establish a HEP community forum (horizontal across theory groups and experiments), linked to broader efforts (e.g., EuCAIF WG1)	MTP
2027-Q1	Phase 2	Benchmark dataset and problem definition	MTP
2028-Q1	Phase 3	Training infrastructure and benchmark suite in place.	MTP
2028-Q4	Phase 3	First models released as shared scientific AI infrastructure.	MTP
2029-Q4	Phase 4	Generalization to realistic dataset	MTP
2030-Q4	Phase 4	Downstream tasks demonstrated, such as end-to-end reconstruction and simulation.	MTP
Horizon 2030+	Horizon 2030+	Continued maturation beyond the plan, subject to resources.	Ext. Funds

5. DEPENDENCIES



Dataset access conditions with the collaborations. Large-scale training capacity from LH-01. Data foundations (CA3).

6. RESOURCES, RISKS AND STATUS

Resources, existing	Resources, new	Top risks, max 3
Ongoing investigations across the HEP community	Effort supported by the MTP allocation, contributions from the R&D WG on Foundation Models, collaborations with external companies	Foundation models might not generalize as expected across different experiments, available datasets might not be representative enough.

Status: Not yet rated **Last update:** 31 August 2026 **Next review:** AIPC plenary, September 2026

CERN AI Agent for Particle Physics

Initiative ID	LH-08	Tier	Horizon 2030+
AIPC Lead	AI for Research Lead	Sector / dept	RCS
Pillar	Pillar 1	Objectives	Obj 1.1, 1.4

1. WHAT IT DELIVERS

Agentic AI capable of assisting physicists in typical tasks related to experimental particle physics and phenomenology, such as performing data analysis, assessing uncertainties, and learning the implications of physics results in terms of a given theoretical result. Starting from the CERN LLM platform, this lighthouse aims at developing a customization pipeline that uses RAG and tuning techniques to evolve the CERN LLM service towards awareness of non-public information, such as experiment internal information for data analysis

2. SCOPE

In scope	Out of scope
Workflow to convert internal information into tokenized datasets for AI agent customization. Workflow to customize an AI agent through post-training techniques (e.g., RAG) on private datasets. Workflow to share knowledge at training time without exposing sensitive data belonging to specific collaborations cross (e.g., through the use of federated learning) Mechanism to prevent information leakage at inference time. Custom AI agents, aware of CERN-internal information and usable for cutting-edge analysis (e.g., LHC data analysis) Software pipeline for LLM customization, exportable to other facilities for non-HEP use cases within the CERN member states	Generic information retrieval (part of LH-02) AI agents for infrastructure (part of LH-01 plus possibly relevant experience in LH-05)

3. EXPECTED IMPACT

By developing such a customization pipeline, this lighthouse will provide the opportunity to produce cutting-edge scientific results by exploiting the demonstrated advantages of commercial LLMs (e.g., Claude and ChatGPT on CERN open data) on the data being collected and analysed by the experiments (e.g., LHC Run 2 and Run 3).

Indicator: To be defined after completion of the exploratory work

4. MILESTONES

Quarter	Phase	Description	Type
2027-Q2	Phase 2	Assess state of the art; identify CERN possible contribution, establish a HEP community forum (horizontal across theory groups and experiments) to discuss developments in the context of ongoing activities	Ext. Funds
2028-Q2	Phase 2	Creation of a private dataset for Agent customization, identification of a model baseline and definition of benchmark tasks covering theoretical and experimental particle physics	Ext. Funds
2028-Q4	Phase 3	Partnership and funding route identified with AI research centres.	Ext. Funds

Quarter	Phase	Description	Type
2029-Q4	Phase 4	Prototype assistant demonstrated on a bounded analysis task, benchmarked against commercial solutions and untuned open models	Ext. Funds
Horizon 2030+	Horizon 2030+	Substantial realisation, subject to additional resources.	Ext. Funds

5. DEPENDENCIES

Additional resources and collaboration with AI research centres. Builds on LH-01 and LH-06.

6. RESOURCES, RISKS AND STATUS

Resources, existing	Resources, new	Top risks, max 3
Several R&D initiatives across the community	External funding and collaborations	Speculative R&D, long-term feasibility and impact to be assessed

Status: Not yet rated **Last update:** 31 August 2026 **Next review:** AIPC plenary, September 2026

AI-Assisted Detector Design

Initiative ID	LH-09	Tier	Horizon 2030+
AIPC Lead	AI for Research Lead	Sector / dept	RCS, ATS
Pillar	Pillar 1	Objectives	Obj 1.1, 1.2

1. WHAT IT DELIVERS

End-to-end optimisation pipeline to assist detector design tasks, treating geometry, materials and technology choices together with the reconstruction and analysis chain as a single objective that balances physics performance against cost, using AI training techniques (e.g., reinforcement learning) and machine-learned surrogate simulators to explore the design space for next-generation detectors, in support of future detector designs (e.g., for HL-LHC Run 5, if dedicated funds and activities are established, and FCC), complementing the fast-simulation and physics-foundation-model initiatives.

2. SCOPE

In scope

Differentiable and surrogate-based optimisation of detector geometry, materials and technology jointly with the reconstruction chain as well as engineering considerations in terms of ease of installation and servicing; benchmarking on reference detector elements such as calorimeter or tracker configurations for HL-LHC upgrades and future-facility concepts. Development of a Reinforcement learning pipeline for problem optimization, deployable for detector design and other use cases, within and beyond HEP

Out of scope

Detector construction, engineering design and procurement; final technology selection by the experiment collaborations; production integration of reconstruction software.

3. EXPECTED IMPACT

A faster, more systematic and reproducible basis for exploring detector-design options against combined physics-performance, engineering and cost objectives, giving future-facility studies a quantitative way to compare concepts ahead of experimental validation. Stated as a direction, not a commitment.

Indicator: A reference detector element optimised end to end and benchmarked against a baseline design.

4. MILESTONES

Quarter	Phase	Description	Type
2027-Q2	Phase 2	Assess state of the art; identify CERN possible contribution; establish link to existing HEP community efforts (e.g., MODE collaboration, EuCAIF WG2) to discuss developments in the context of ongoing activities	Ext. Funds
2028-Q4	Phase 3	AI-driven and surrogate-based optimisation assessed on a reference detector element.	Ext. Funds
2029-Q4	Phase 4	Collaboration and funding route identified with detector R&D programmes and the AI-for-detector-design community.	Ext. Funds
Horizon 2030+	Horizon 2030+	Standardization of the procedure into a consolidated software library covering all use cases (tracking, calorimetry, PID, etc.), subject to additional resources and collaborations.	Ext. Funds

5. DEPENDENCIES

Additional resources and collaborations. Alignment with future-detector studies for HL-LHC upgrades and the FCC. Builds on LH-04 (fast simulation), LH-06 (AI-assisted accelerator systems design), and LH-07 (physics foundation model).

6. RESOURCES, RISKS AND STATUS

Resources, existing	Resources, new	Top risks, max 3
Several R&D initiatives in the community	External resources, grants, collaborations	Speculative R&D, long-term feasibility and impact to be assessed

Status: Not yet rated **Last update:** 31 August 2026 **Next review:** AIPC plenary, September 2026

AI-Driven Material Discovery for Future Colliders

Initiative ID	LH-10	Tier	Horizon 2030+
AIPC Lead	AI for Research and AI for Accelerators Leads	Sector / dept	ATS
Pillar	Pillar 1	Objectives	Obj 1.1, 1.2

1. WHAT IT DELIVERS

AI-driven discovery and inverse design of advanced materials for future-colliders and experiment technologies, such as high-field superconductors and radiation-hard materials, using generative models and machine-learned property prediction to narrow the candidate space ahead of experimental validation, in support of FCC magnet and conductor R&D.

2. SCOPE

In scope

AI-based HTS material discovery study for FCC accelerator and experimental magnets
Development of AI-based techniques for magnet design
AI-based optimization of the industrialization procedure

Out of scope

Prototype production
Demonstrator deployment

3. EXPECTED IMPACT

Building on the impact AI has had in materials discovery and chemistry more broadly (for example, the AlphaFold work by DeepMind, recognised by the 2024 Nobel Prize in Chemistry), this lighthouse aims to develop a CERN-led consortium of experts to accelerate HTS research with AI techniques. A successful deliverable of HTS material could have a wide variety of positive impacts on non-HEP applications, as well as providing already concrete solutions for FCCee and fostering research towards FCChh magnet design, that could help CERN meet the goals of the FCChh timeline.

Indicator: To be defined after completion of the exploratory work

4. MILESTONES

Quarter	Phase	Description	Type
2027-Q4	Phase 2	Preparation started, candidate material classes and validation route defined with FCC magnet and conductor R&D.	Ext. Funds
2028-Q4	Phase 3	Generative models and property prediction assessed on a reference material class.	Ext. Funds
2029-Q4	Phase 4	Collaboration and funding route identified.	Ext. Funds
Horizon 2030+	Horizon 2030+	Substantial realisation, subject to additional resources and collaborations.	Ext. Funds

5. DEPENDENCIES

Additional resources and collaborations. Alignment with FCC magnet and conductor R&D programme.

6. RESOURCES, RISKS AND STATUS

Resources, existing	Resources, new	Top risks, max 3
Several R&D initiatives in the community and industry	External resources, grants, collaborations	Speculative R&D, long-term feasibility and impact to be assessed

Status: Not yet rated **Last update:** 31 August 2026 **Next review:** AIPC plenary, September 2026



CERN
CH-1211 Geneva, Switzerland