ETP4HPC'S SRA 7



Strategic Research Agenda for High-Performance Computing in Europe

European HPC Research Priorities for 2028 - 2031

NOVEMBER 2026

ETP4HPC.EU

Table of contents

Message from the ETP4HPC's Chairman 5
Glossary of Acronyms 6
Use of Generative AI 11
1 Executive Summary of Recommendations 12
1.1 European Sovereignty, Industrial Adoption and Ecosystem Leadership 13
R1 - Strengthen European AI Sovereignty and Accelerate AI-Driven Scientific Discovery 13
R2 - Make Post-Exascale a coordinated European mission 13
R3 - Give European SMEs and industry easier access to Advanced Computing 14
R4 - Coordinate skills, governance and long-term ecosystem stewardship 14
1.2 Federated, Sustainable and Trustworthy Computing Infrastructure 14
R5 - Build a genuinely federated European computing continuum 14
R6 – Establish Enforceable Interoperability Requirements in European HPC Procurement 14
R7 - Put energy efficiency and sustainability into every investment decision 15
R8 - Make software provenance tracing and cyber-resilience mandatory 15
1.3 European Computing Technologies and Architectures 15
R9 - Continue Investments in European processors, accelerators and interconnects 15
R10 - Establish a European chiplet and advanced-packaging ecosystem 15
R11 - Create a modular experimental facility for disruptive computing 16
R12 - Advance Open, Secure and Scalable Quantum–HPC Integration 16
1.4 Software, Algorithms and Intelligent Computing 16
R13 - Support technologies and methods to optimize data sharing and data access 16
R14 - Fund sovereign European "software factories" 16
R15 - Modernise resource management for dynamic workflows 16
R16 – Advance Scalable, Energy-Efficient and Performance-Portable Software and Algorithms for Post-Exascale Computing 17
2 SRA7 Introduction and Structure 20
2.1 So why does the SRA exist and what are the primary objectives in SRA7? 21
2.1.1 SRA7, a community perspective 21
2.1.2 SRA7, a policy makers' perspective 22
2.1.3 SRA7 structure and release 22
2.1.4 SRA7 white papers and master 23
2.2 Continued strong linkages with EuroHPC and the EC 24
2.3 Continued strong engagement with the ecosystem communities 24
2.3.1 Strong focus on community feedback 25
3 European Advanced Computing - Overall State of Play 26
3.1 AI Factories and the rapid expansion of AI infrastructure 27

3.2 Europe enters the exascale era 28
3.3 Quantum computing becomes an operational pillar 28
3.4 European technological sovereignty and the RISC-V ecosystem 30
3.5 Federation of computing, data and software resources 31
3.6 Sustainability, software and applications 31
3.7 Skills, widening and international cooperation 31
3.8 Conclusion 33
4 Advanced Computing: R&I Recommendations 34
4.1 Research Domain Recommendations 35
4.1.1 System Architecture 35
4.1.2 System Hardware Components 36
4.1.3 System Software and Management 38
4.1.4 Programming Environments 41
4.1.5 Data Management I/O and Storage 43
4.1.6 Mathematics and Algorithms 44
4.1.7 European Computing Continuum 47
4.2 Cross cutting technical theme recommendations 49
4.2.1 Non-Conventional Systems and Architectures 49
4.2.2 Quantum Computing 50
4.2.3 HPC/AI Usage, Applications and Workflows 52
4.2.4 Energy Efficiency and Sustainability 54
5 Industrial Usage: Priorities for European HPC, AI and Data Ecosystems 56
5.1 Industrial context and strategic objective 57
5.2 Strategic research and innovation priorities 57
5.2.1 Sovereign industrial AI and HPC convergence 57
5.2.2 Industrial digital twins and real-time simulation 58
5.2.3 Emerging computing architectures for industrial workloads 58
5.2.4 Quantum readiness for industrial applications 59
5.2.5 Trusted, secure and sovereign industrial computing 59
5.2.6 Sustainable and energy-aware HPC and AI 59
5.2.7 Industrial workflow automation and autonomous engineering 60
5.2.8 A federated HPC–cloud–edge continuum 60
5.2.9 Portable software and resilient software ecosystems 60
5.2.10 Industrial data spaces and semantic interoperability 61
5.2.11 Industrial skills, adoption and competence ecosystems 61
5.3 Measuring industrial impact 61
6 European HPC Technology SMEs 62
6.1 State of play 63

6.2 Strategic recommendations 63
6.2.1 Recognise HPC technology SMEs as a strategic capability class 63
6.2.2 Create SME-first tracks in European programmes and initiatives 63
6.2.3 Reform procurement to move SMEs from subcontractors to prime contractors 63
6.2.4 Bridge the prototype-to-product gap for European HPC technology SMEs 64
6.2.5 Build sovereign AI/HPC software and middleware around SMEs 64
6.2.6 Make AI Factories and Gigafactories SME product engines 64
6.2.7 Establish a European HPC SME catalogue and coordination hub 65
6.2.8 Create SME-integrated talent pathways 65
6.2.9 Establish an SME policy and political advocacy group 65
6.2.10 Measure what matters 65
6.3 Conclusion 65
7 Conclusion and Acknowledgements 68
7.1 Acknowledgements 70

Message from the ETP4HPC's Chairman



Two years is a long time in Advanced Computing. When SRA6 was published, Europe was still waiting for its first exascale machine; today JUPITER is running, Alice Recoque is on its way, and the boundaries between HPC, AI and quantum computing have largely faded. The EuroHPC Joint Undertaking has grown accordingly. Its mandate now explicitly covers AI, nineteen AI Factories are coming online, the Gigafactories call is open, and European quantum computers have welcomed their first users. With the Federation Platform, these assets are beginning to work as one continuum, and the JU now reaches across the Atlantic, Canada having become its 39th participating state.

This progress should not hide an uncomfortable fact. The first European processor is a real achievement, yet most of the components, software and AI models running in our machines are still designed elsewhere. As geopolitical tensions rise, this dependency is no longer only an industrial concern; it touches Europe's security and its freedom to choose its own path.

SRA7 is our answer. It sets out what Europe should research and build over the next five years: sovereign processors and interconnects, a chiplet ecosystem, a truly federated continuum, shared software factories, and energy efficiency as a criterion in every investment. Its timing matters, as the negotiations on the next Multiannual Financial Framework, the European Competitiveness Fund and the next Framework Programme will shape Europe's computing ambitions for the decade to come.

This agenda is the work of some 130 experts from industry, research and academia, guided by our working group leaders and brought together by the editorial team. On behalf of ETP4HPC, I thank them warmly. Thirteen white papers will follow, then agile papers through 2027 and early 2028. Our hope is simple: that these recommendations shape the research and innovation actions of the years ahead.

Dominique Grelet

Chair, ETP4HPC

Glossary of Acronyms

Acronyms and abbreviations used throughout this document are defined in the Glossary.

AC — Advanced Computing; computing technologies and systems encompassing HPC, AI, quantum computing and related domains.

ABFT — Algorithm-Based Fault Tolerance; techniques for detecting or correcting computational errors using algorithmic properties.

AI — Artificial Intelligence; computational methods and systems capable of performing tasks associated with human intelligence.

API — Application Programming Interface; defined interface enabling communication between software components.

BXI — Bull eXascale Interconnect; high-performance interconnect technology developed by Bull for HPC systems.

Carbon-to-Solution — Greenhouse-gas emissions attributable to completing a defined computational workload, typically expressed in CO_2-equivalent.

CASTIEL — European coordination and support initiative connecting national HPC competence centres.

CCI — Computing Continuum Initiative; collaboration addressing computing across HPC, cloud, edge and related environments.

CERN — European Organization for Nuclear Research; European research organisation operating large-scale scientific infrastructure.

CI/CD — Continuous Integration and Continuous Delivery/Deployment; automated software build, test and release practices.

cPPP — Contractual Public–Private Partnership; cooperation instrument used under Horizon 2020.

CPU — Central Processing Unit; general-purpose processor executing program instructions.

CSRD — Corporate Sustainability Reporting Directive; European Union sustainability reporting legislation.

CUE — Carbon Usage Effectiveness; carbon emissions relative to IT equipment energy consumption.

CVE — Common Vulnerabilities and Exposures; standard identifiers for publicly disclosed cybersecurity vulnerabilities.

CycloneDX — Software supply-chain bill-of-materials standard for representing components, dependencies and related security information.

DARE — Digital Autonomy with RISC-V in Europe; project advancing European processor and accelerator technologies.

DC-SCM — Data Center Secure Control Module; modular security and management architecture for data-centre platforms.

DDR — Double Data Rate; family of synchronous DRAM memory technologies.

DG CNECT — Directorate-General for Communications Networks, Content and Technology; European Commission department responsible for digital policies.

DMR — Dual Modular Redundancy; duplicated execution for fault detection.

DPU — Data Processing Unit; programmable processor designed to accelerate data movement and infrastructure services, including networking, storage and security.

DSL – Domain Specific Language

EC — European Commission; executive institution of the European Union.

EDA — Electronic Design Automation; software tools used to design and verify electronic systems.

EED — Energy Efficiency Directive; European Union legislation concerning energy efficiency.

EESSI — European Environment for Scientific Software Installations; shared software distribution infrastructure for scientific computing.

EMECF — European Modular Experimental Computing Facility; proposed modular infrastructure for evaluating non-conventional computing technologies.

Energy-to-Solution — Total energy consumed to complete a defined computational workload; metric for evaluating application-level energy efficiency.

EOSC — European Open Science Cloud; European framework for federated research data and services.

EPI — European Processor Initiative; European programme developing processor technologies for HPC.

ERF — Energy Reuse Factor; fraction of facility energy that is reused.

ETP4HPC — European Technology Platform for High Performance Computing; European industry-led association supporting HPC research and innovation.

EU — European Union; political and economic union of European member states.

EUMaster4HPC — European Master for High Performance Computing; European education programme developing advanced HPC expertise.

EuroCC — European network of National Competence Centres supporting HPC adoption, skills and access.

EuroHPC JU — European High Performance Computing Joint Undertaking; European partnership supporting HPC, AI and quantum computing infrastructure and research.

EUPEX — European Pilot for Exascale; project developing and validating a European HPC software and hardware ecosystem.

FAIR — Findable, Accessible, Interoperable and Reusable; principles for digital research resources.

FP32 — 32-bit Floating Point; single-precision numerical format.

FP64 — 64-bit Floating Point; double-precision numerical format widely used in scientific computing.

FPGA — Field-Programmable Gate Array; reconfigurable integrated circuit that can be programmed to implement application-specific hardware functions and accelerators.

FTQC — Fault-Tolerant Quantum Computing; quantum computing using error correction to support reliable logical operations.

GDPR — General Data Protection Regulation; European Union regulation governing personal data protection.

GitOps — Git-based Operations; approach to infrastructure management using version-controlled, declarative configurations.

GPU — Graphics Processing Unit; highly parallel processor used for graphics, AI and scientific computing.

HBM — High Bandwidth Memory; stacked memory providing high data-transfer bandwidth.

HiPEAC — High Performance, Edge And Cloud computing; European computing research and innovation network.

HPCG — High Performance Conjugate Gradient; benchmark measuring computational and memory-system behaviour.

HPCQS — High Performance Computer and Quantum Simulator Hybrid; European project integrating quantum simulators with HPC infrastructure.

HPCTRAIN — European HPC professional traineeship programme supporting skills development and professional training.

IEC — International Electrotechnical Commission; international standards organisation for electrical and electronic technologies.

IEEE — Institute of Electrical and Electronics Engineers; international professional organisation developing technical standards.

INFRAG — Infrastructure Advisory Group; advisory group supporting EuroHPC JU infrastructure strategy.

IP — Intellectual Property; legally protected inventions, designs, software and other intellectual creations.

IPU — Infrastructure Processing Unit; specialised processor designed to offload and accelerate infrastructure functions such as networking, storage, security and virtualisation.

ITER — International Thermonuclear Experimental Reactor; international fusion-energy research project.

JU — Joint Undertaking; European partnership established to implement strategic research and innovation programmes.

LCA — Life Cycle Assessment; methodology for evaluating environmental impacts across the lifecycle of a product or system, including manufacturing, operation and end-of-life.

LLM — Large Language Model; AI model trained on large datasets to process and generate language.

MASP — Multi-Annual Strategic Programme; EuroHPC Joint Undertaking's multiannual strategic planning document.

ML-DSA — Module-Lattice-Based Digital Signature Algorithm; post-quantum digital signature standard.

ML-KEM — Module-Lattice-Based Key-Encapsulation Mechanism; post-quantum cryptographic standard for establishing shared secrets.

MLOps — Machine Learning Operations; practices for deploying, monitoring and maintaining machine-learning systems.

MPI — Message Passing Interface; standard for communication between processes in parallel computing.

MRDIMM — Multiplexed Rank Dual In-line Memory Module; high-bandwidth server memory module technology.

NCC(s) — National Competence Centre(s); national centres supporting HPC adoption, training, skills development and access to HPC expertise and services.

NET4EXA — European research project developing high-performance interconnect technologies for exascale computing.

NIC — Network Interface Card; hardware component providing connectivity between a computing node and a network.

NIST — National Institute of Standards and Technology; US standards and technology research organisation.

NoC — Network-on-Chip; on-chip communication architecture connecting processing elements, memory controllers and other components within a system-on-chip or accelerator.

NPU — Neural Processing Unit; specialised processor designed to accelerate artificial intelligence and neural-network workloads.

NVLink — NVIDIA NVLink; high-bandwidth, low-latency interconnect technology enabling communication between GPUs and, in supported architectures, between CPUs and GPUs.

PIM — Processing-in-Memory; computation performed within or near memory to reduce data transfers.

PMIx — Process Management Interface for Exascale; standard interfaces supporting parallel process and resource management.

PUE — Power Usage Effectiveness; total facility energy divided by IT equipment energy.

Q-CAB — Quantum Computing Advisory Board; advisory body established by the European Union.

QDMI — Quantum Device Management Interface; interface for accessing and managing quantum devices.

QEC — Quantum Error Correction; techniques for protecting quantum information against errors and decoherence.

Q-Neko — EU–Japan research cooperation project on hybrid quantum–HPC algorithms.

QPU — Quantum Processing Unit; hardware that executes quantum operations.

QRMI — Quantum Resource Management Interface; interface concepts for managing quantum resources.

QTAG — Quantum Technologies Advisory Group; advisory group supporting EuroHPC JU's quantum computing activities.

RIAG — Research and Innovation Advisory Group; advisory group of the EuroHPC Joint Undertaking.

RISC-V — Reduced Instruction Set Computer – Five; open-standard instruction-set architecture.

RJMS — Resource and Job Management System; software for allocating computing resources and scheduling jobs.

RO-Crate — Research Object Crate; packaging approach for research data and structured metadata to support discovery, reuse and reproducibility.

SBOM — Software Bill of Materials; machine-readable inventory of software components and dependencies.

SDC — Silent Data Corruption; undetected corruption of data or computational results.

SLH-DSA — Stateless Hash-Based Digital Signature Algorithm; post-quantum digital signature standard.

SLSA — Supply-chain Levels for Software Artifacts; framework for improving software supply-chain integrity.

SmartNIC — Smart Network Interface Card; programmable network adapter capable of offloading networking and infrastructure processing.

SME — Small and Medium-sized Enterprise; business category defined by employee count and financial thresholds.

SPDX — Software Package Data Exchange; open standard for communicating software components, licences and provenance information.

SPECpower — SPECpower_ssj2008; benchmark for server performance and power consumption.

SPE-EU — European project supporting post-exascale roadmapping and strategic coordination.

SRA — Strategic Research Agenda; document identifying research and innovation priorities.

SSO — Single Sign-On; authentication mechanism allowing users to access multiple services with one login.

TCI — TransContinuum Initiative; European collaboration addressing the computing continuum.

TMR — Triple Modular Redundancy; triplicated execution with majority voting for fault tolerance.

TRL — Technology Readiness Level; scale for assessing the maturity of a technology.

UALink — Ultra Accelerator Link; open interconnect technology for communication between accelerators.

UCIe — Universal Chiplet Interconnect Express; open die-to-die interconnect standard.

UEC — Ultra Ethernet Consortium; industry consortium developing Ethernet-based networking specifications for AI and HPC.

UPS — Uninterruptible Power Supply; system providing temporary power during supply interruptions.

W3C PROV — World Wide Web Consortium Provenance Data Model and related specifications; standards for representing and exchanging provenance information.

WUE — Water Usage Effectiveness; water consumption per unit of IT energy.

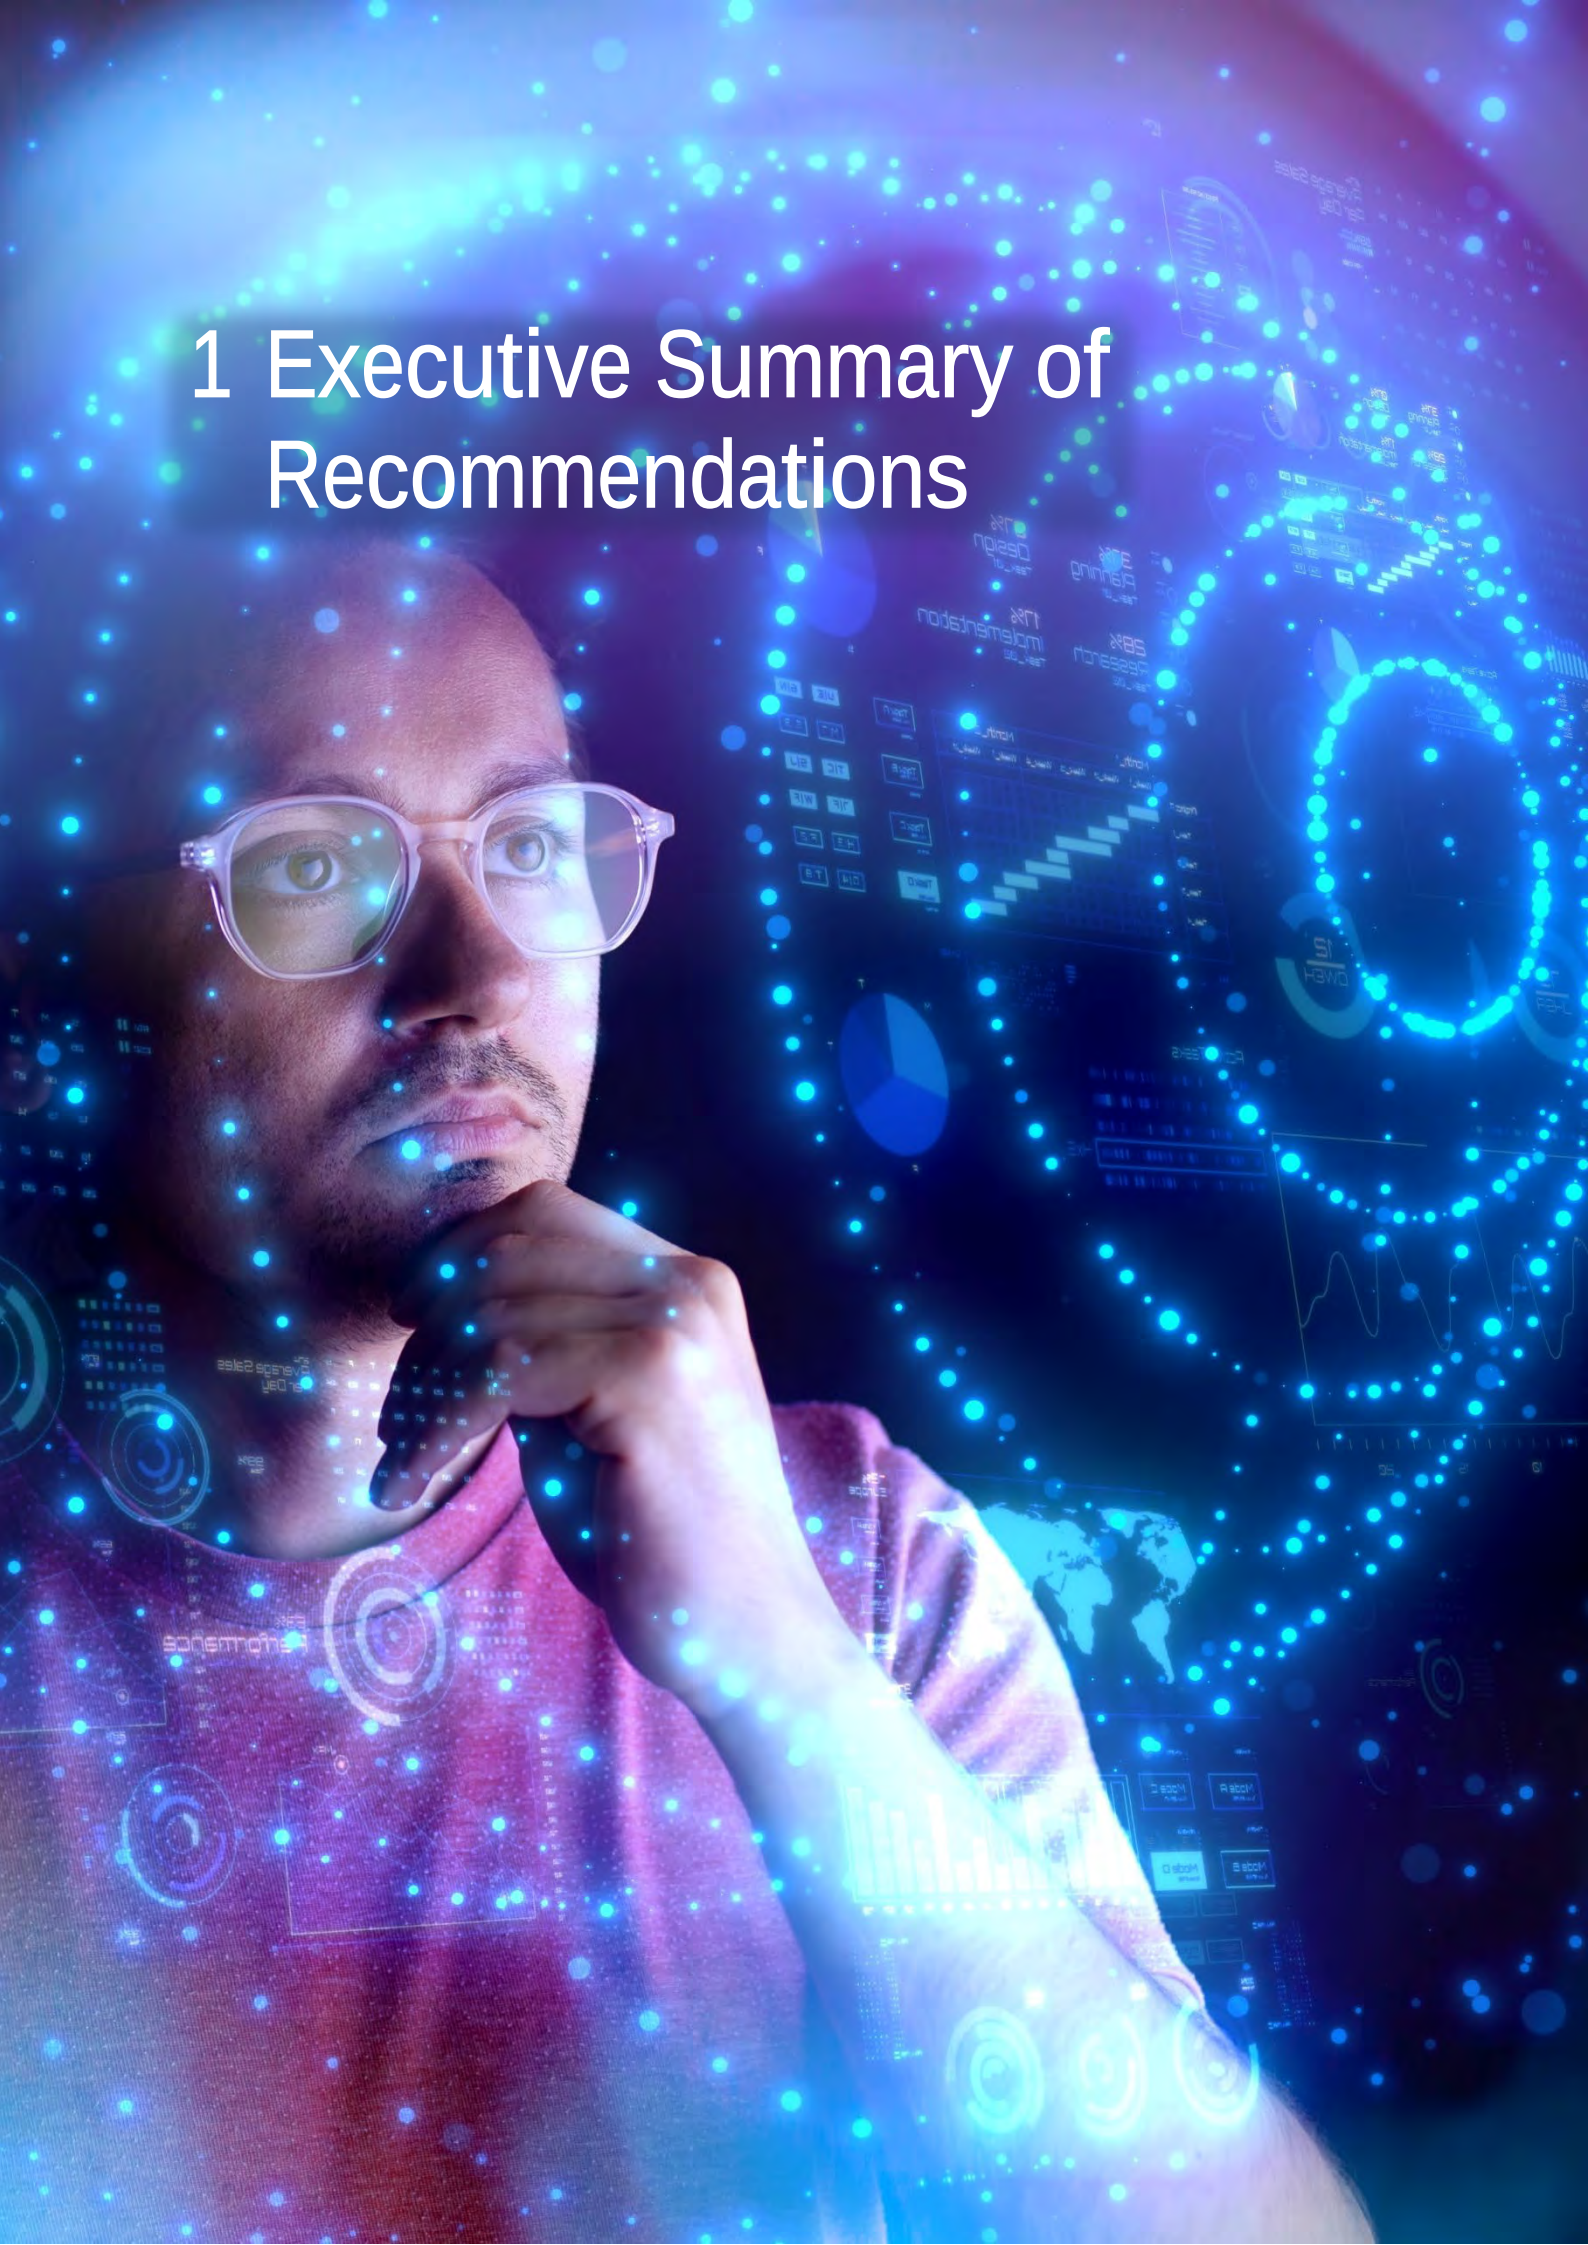
Use of Generative AI

Generative Artificial Intelligence (GenAI) tools were used in the preparation of selected parts of this document as editorial support for drafting, reviewing and refining the text, including improvements to clarity, structure, consistency and language.

The substantive content, technical insights, research priorities and strategic recommendations presented in this Strategic Research Agenda originate from the expertise, contributions and collective deliberations of the ETP4HPC working groups, contributing experts and the editorial team. GenAI was not used as an authoritative source of scientific knowledge or as a substitute for expert judgement in defining the core content and strategic recommendations.

The final content and recommendations remain the responsibility of the ETP4HPC (SRA7 editorial team and the SRA7 White Paper authors).

1 Executive Summary of Recommendations



The following are the top-level recommendations coming out of the 2026 edition of the ETP4HPC Strategic Research Agenda. More details are in the SRA7 White Papers and there is more context provided in the following chapters of this report. All these recommendations are primarily targeted towards EC, EuroHPC and the Member States to be actioned in the 2027 – 2032 timeframe. [1] At the end we also provide a table mapping these Recommendations to closest matching white papers. We show here the recommendations across the following themes:

- European Sovereignty, Industrial Adoption and Ecosystem Leadership
- Federated, Sustainable and Trustworthy Computing Infrastructure
- European Computing Technologies and Architectures
- Software, Algorithms and Intelligent Computing

1.1 European Sovereignty, Industrial Adoption and Ecosystem Leadership

R1 - Strengthen European AI Sovereignty and Accelerate AI-Driven Scientific Discovery

Europe should continue to place greater emphasis on developing a **coordinated strategy for AI in Advanced Computing** that addresses critical dependencies and supply-chain risks across AI hardware, foundational models, training datasets, software frameworks and infrastructure, while strengthening European technological sovereignty. Research and Innovation investments should prioritise the development and **adoption of trustworthy, open and sovereign AI models tailored to HPC and scientific applications**, including domain-specific foundation models, AI-enhanced simulations, surrogate models and hybrid AI–HPC workflows. Europe should also support the development of **AI-powered scientific assistants and autonomous research agents** to support scientists and industrial innovators in hypothesis generation, experiment design, simulation orchestration, data analysis and scientific discovery – leading to faster innovation throughput and societal breakthroughs - drawing inspiration from other similar initiatives globally [2]. Europe should also support the development of open **European benchmarks and datasets for AI-assisted science.**

R2 - Make Post-Exascale a coordinated European mission

Europe should develop a co-ordinated multi-year **post-exascale research and innovation roadmap** with very well-defined milestones for architecture, software, applications, energy and deployment, in HPC, AI and Quantum. There is a need to align EuroHPC, Horizon Europe and relevant Chips JU actions around measurable **chip-level and system-level outcomes for Europe** rather than isolated technology demonstrations. Projects such as SPE-EU [3] are a good first step in this direction but more is needed in terms of alignment across different community and JU roadmaps. Such an exercise will help to identify

[1] The executive summary lists top level priorities in the various areas of Advanced Computing, so the ordering of these do not have any bearing on the priorities.
[2] https://www.energy.gov/undersecretaryforscience/genesis-mission/genesis-mission
[3] https://spe-project.eu/

and promote disruptive new computing paradigms for the post-exascale era where Europe can start to gain strategic leadership.

R3 - Give European SMEs and industry easier access to Advanced Computing

Europe should expand **accessible services (Incl. AI Factories and Gigafactories), support centres, testbeds and procurement pathways for SMEs and industrial users**. Reducing onboarding and specialist-skills barriers through unified portals, training and transparent access models is necessary. These measures are essential to democratising access to Advanced Computing and AI capabilities, enabling European SMEs and industries to innovate, compete globally and translate Europe's investments in digital infrastructure into sustainable economic growth and technological leadership.

R4 - Coordinate skills, governance and long-term ecosystem stewardship

Europe should establish a coordinated strategy involving EuroHPC, the European Commission, research organisations and industry to develop and **retain world-class talent in HPC, AI and quantum computing, with particular emphasis on cross-disciplinary skills**, international mobility and long-term career opportunities. As competition for highly skilled professionals intensifies, particularly from non-European hyperscalers and AI model providers, Europe must strengthen its ability to attract and retain talent while systematically evaluating the effectiveness and impact of existing initiatives[4], to ensure that investments in skills translate into sustained European scientific leadership, industrial competitiveness and technological sovereignty. Higher education and specialised training programmes/curricula to support the next generation of experts in the design, deployment, operation, and effective use of current and future HPC and Advanced Computing technologies should be developed. These programmes should prioritise interdisciplinary skills, practical experience, and industry–academia collaboration.

1.2 Federated, Sustainable and Trustworthy Computing Infrastructure

R5 - Build a genuinely federated European computing continuum

Europe should prioritise **production-ready federation** of HPC, AI Factories, quantum resources, cloud, edge and data infrastructures moving forward with bolder ambition with some of the initiatives that have already started[5]. There is a need to fund interoperable identity, resource discovery, allocation, workflow execution and accounting across national boundaries, with multi-site pilots.

R6 - Establish Enforceable Interoperability Requirements in European HPC Procurement

Europe should move beyond encouraging interoperability towards **making open interfaces, software portability and demonstrable cross-vendor interoperability explicit requirements in publicly funded HPC and AI procurements**, where technically appropriate. The European Commission, EuroHPC JU and national procuring authorities should incorporate these requirements into procurement technical specifications, award criteria and contractual acceptance tests, consistent with applicable EU procurement legislation. Conformance testing should cover accelerators, networks, storage, orchestration and federated services, using open standards and verifiable interfaces wherever

[4] Such as https://eumaster4hpc.eu/
[5] https://my-eurohpc.eu/

available, to reduce vendor lock-in and preserve European technological freedom of choice. Proportionate exemptions should be permitted for research prototypes, experimental architectures and first-of-a-kind pilot systems, where strict interoperability requirements could constrain innovation. Such exemptions should be justified, documented and accompanied by a roadmap towards interoperability, where relevant.

R7 - Put energy efficiency and sustainability into every investment decision

Europe should introduce **energy, carbon, water and lifecycle metrics** across infrastructure procurement, operations and research evaluations. **Energy-aware scheduling, waste-heat reuse and AI-assisted tuning**, with cross-site telemetry and independently verifiable results should be supported. This is much more critical today with increasing power consumption expected from European Advanced Computing infrastructures and major geopolitical constraints in energy supply, and with the need to adhere to the European Green Deal.

R8 - Make software provenance tracing and cyber-resilience mandatory

Europe should establish common security, transparency and reliability requirements for publicly funded HPC and AI software and infrastructure. Research and innovation investments should prioritise secure software development, trustworthy supply chains, traceability and protection against emerging cyber threats, including those associated with quantum computing. There should be a requirement to produce **software bills of materials, signed build attestations, vulnerability scanning and reproducibility** for publicly supported software distributions. There is also a need to embed **data lineage, security-by-design and post-quantum migration** planning into federated infrastructure and research workflows.

1.3 European Computing Technologies and Architectures

R9 - Continue Investments in European processors, accelerators and interconnects

Europe should continue to sustain **long-term hardware programmes such as RISC-V-based compute, specialised accelerators and high-bandwidth low-latency networking**. Developing and nurturing market readiness towards a fully working European Advanced Computing hardware ecosystem will take time and patience, with some very good activities that are already underway starting with the availability of the European Processor[6]. Coupling silicon development with compiler, runtime, library and application ecosystems, and creating viable pathways also for European SMEs in this area to reach production deployments is essential.

R10 - Establish a European chiplet and advanced-packaging ecosystem

Europe should support interoperable chiplets, active and photonic interposers, memory integration and advanced packaging. Developing shared interfaces, validation facilities and supply-chain partnerships that allow **European intellectual property in this area to be integrated economically into future HPC and AI systems is required**. Without these capabilities, Europe risks becoming only a supplier of semiconductor intellectual property rather than a competitive developer of complete HPC

[6] https://www.innovationnewsnetwork.com/european-processor-initiative-project-concluded/72029/

and AI systems, with much of the economic value, system-level innovation and strategic control captured outside Europe.

R11 - Create a modular experimental facility for disruptive computing

Europe should assess the case for and finance a **"European Modular Experimental Computing Facility" (EMECF)** to validate non-conventional architectures, software stacks and algorithms at a realistic scale. Structured access to start-ups, SMEs and researchers, with transparent technical assessment and routes to industrialisation should be offered. It would help bridge the gap between research prototypes and industrial deployment by enabling systematic evaluation against realistic workloads, identifying technical barriers and assessing readiness for integration into production HPC and AI environments.

R12 - Advance Open, Secure and Scalable Quantum–HPC Integration

Europe should invest in open standards, interoperable architectures and secure computing environments that enable the seamless integration of quantum and classical HPC systems, while avoiding vendor lock-in and fostering European innovation. Research and innovation should prioritise **scalable quantum error correction, secure multi-user access, and common frameworks for resource estimation and performance benchmarking** to accelerate the transition towards **fault-tolerant quantum computing (FTQC)**, demonstrate measurable benefits and guide future research and industrial investments. This is an opportunity for Europe to gain and show leadership in this topic, and it should seize this!

1.4 Software, Algorithms and Intelligent Computing

R13 - Support technologies and methods to optimize data sharing and data access

Europe should prioritise intelligent, efficient and interoperable data management technologies enabling seamless access to distributed data and supporting high-performance data transfers across computing infrastructures, while reducing energy consumption and operational costs. There is a need to fund solutions, data management methodologies and technologies that support bringing compute-to-data, doing **in-situ and in-transit processing, and performing locality-aware orchestration and intelligent data placement**.

R14 - Fund sovereign European "software factories"

Europe should invest in shared, sustainable infrastructure for the development, integration, testing and deployment of HPC and AI software, ensuring interoperability, reliability and portability across European computing systems. **Shared, continuously operated integration, build, test and deployment services for European HPC and AI software should be built**. There is a need to adopt reusable packaging and deployment approaches, supporting maintainers over the long term and validating stacks across multiple architectures and EuroHPC sites.

R15 - Modernise resource management for dynamic workflows

Europe should invest in flexible, intelligent and interoperable resource management technologies that enable efficient and seamless execution of HPC and AI workloads across diverse computing infrastructures, while ensuring reliability, transparency and human oversight. Methods and techniques

in **malleable jobs, dynamic resource pools, container-native execution and workflow-aware scheduling should be further developed**. Demonstration of **portable orchestration across HPC and AI workloads**, and enablement of human-supervised AI assistance without sacrificing reliability or auditability is needed.

R16 – Advance Scalable, Energy-Efficient and Performance-Portable Software and Algorithms for Post-Exascale Computing

Europe should prioritise Research and Innovation investments in next-generation algorithms, mathematical methods, programming models and software environments that efficiently exploit heterogeneous HPC, AI and quantum computing architectures. Emphasis should be placed on **performance portability, scalability, energy efficiency, numerical accuracy and resilience**, including adaptive precision, communication-avoiding algorithms, fault tolerance and AI-assisted software development. Europe should support open, interoperable programming interfaces and numerical libraries, ensuring long-term software sustainability, reducing vendor lock-in and enabling efficient execution across diverse computing platforms.

1. Executive Summary of Recommendations

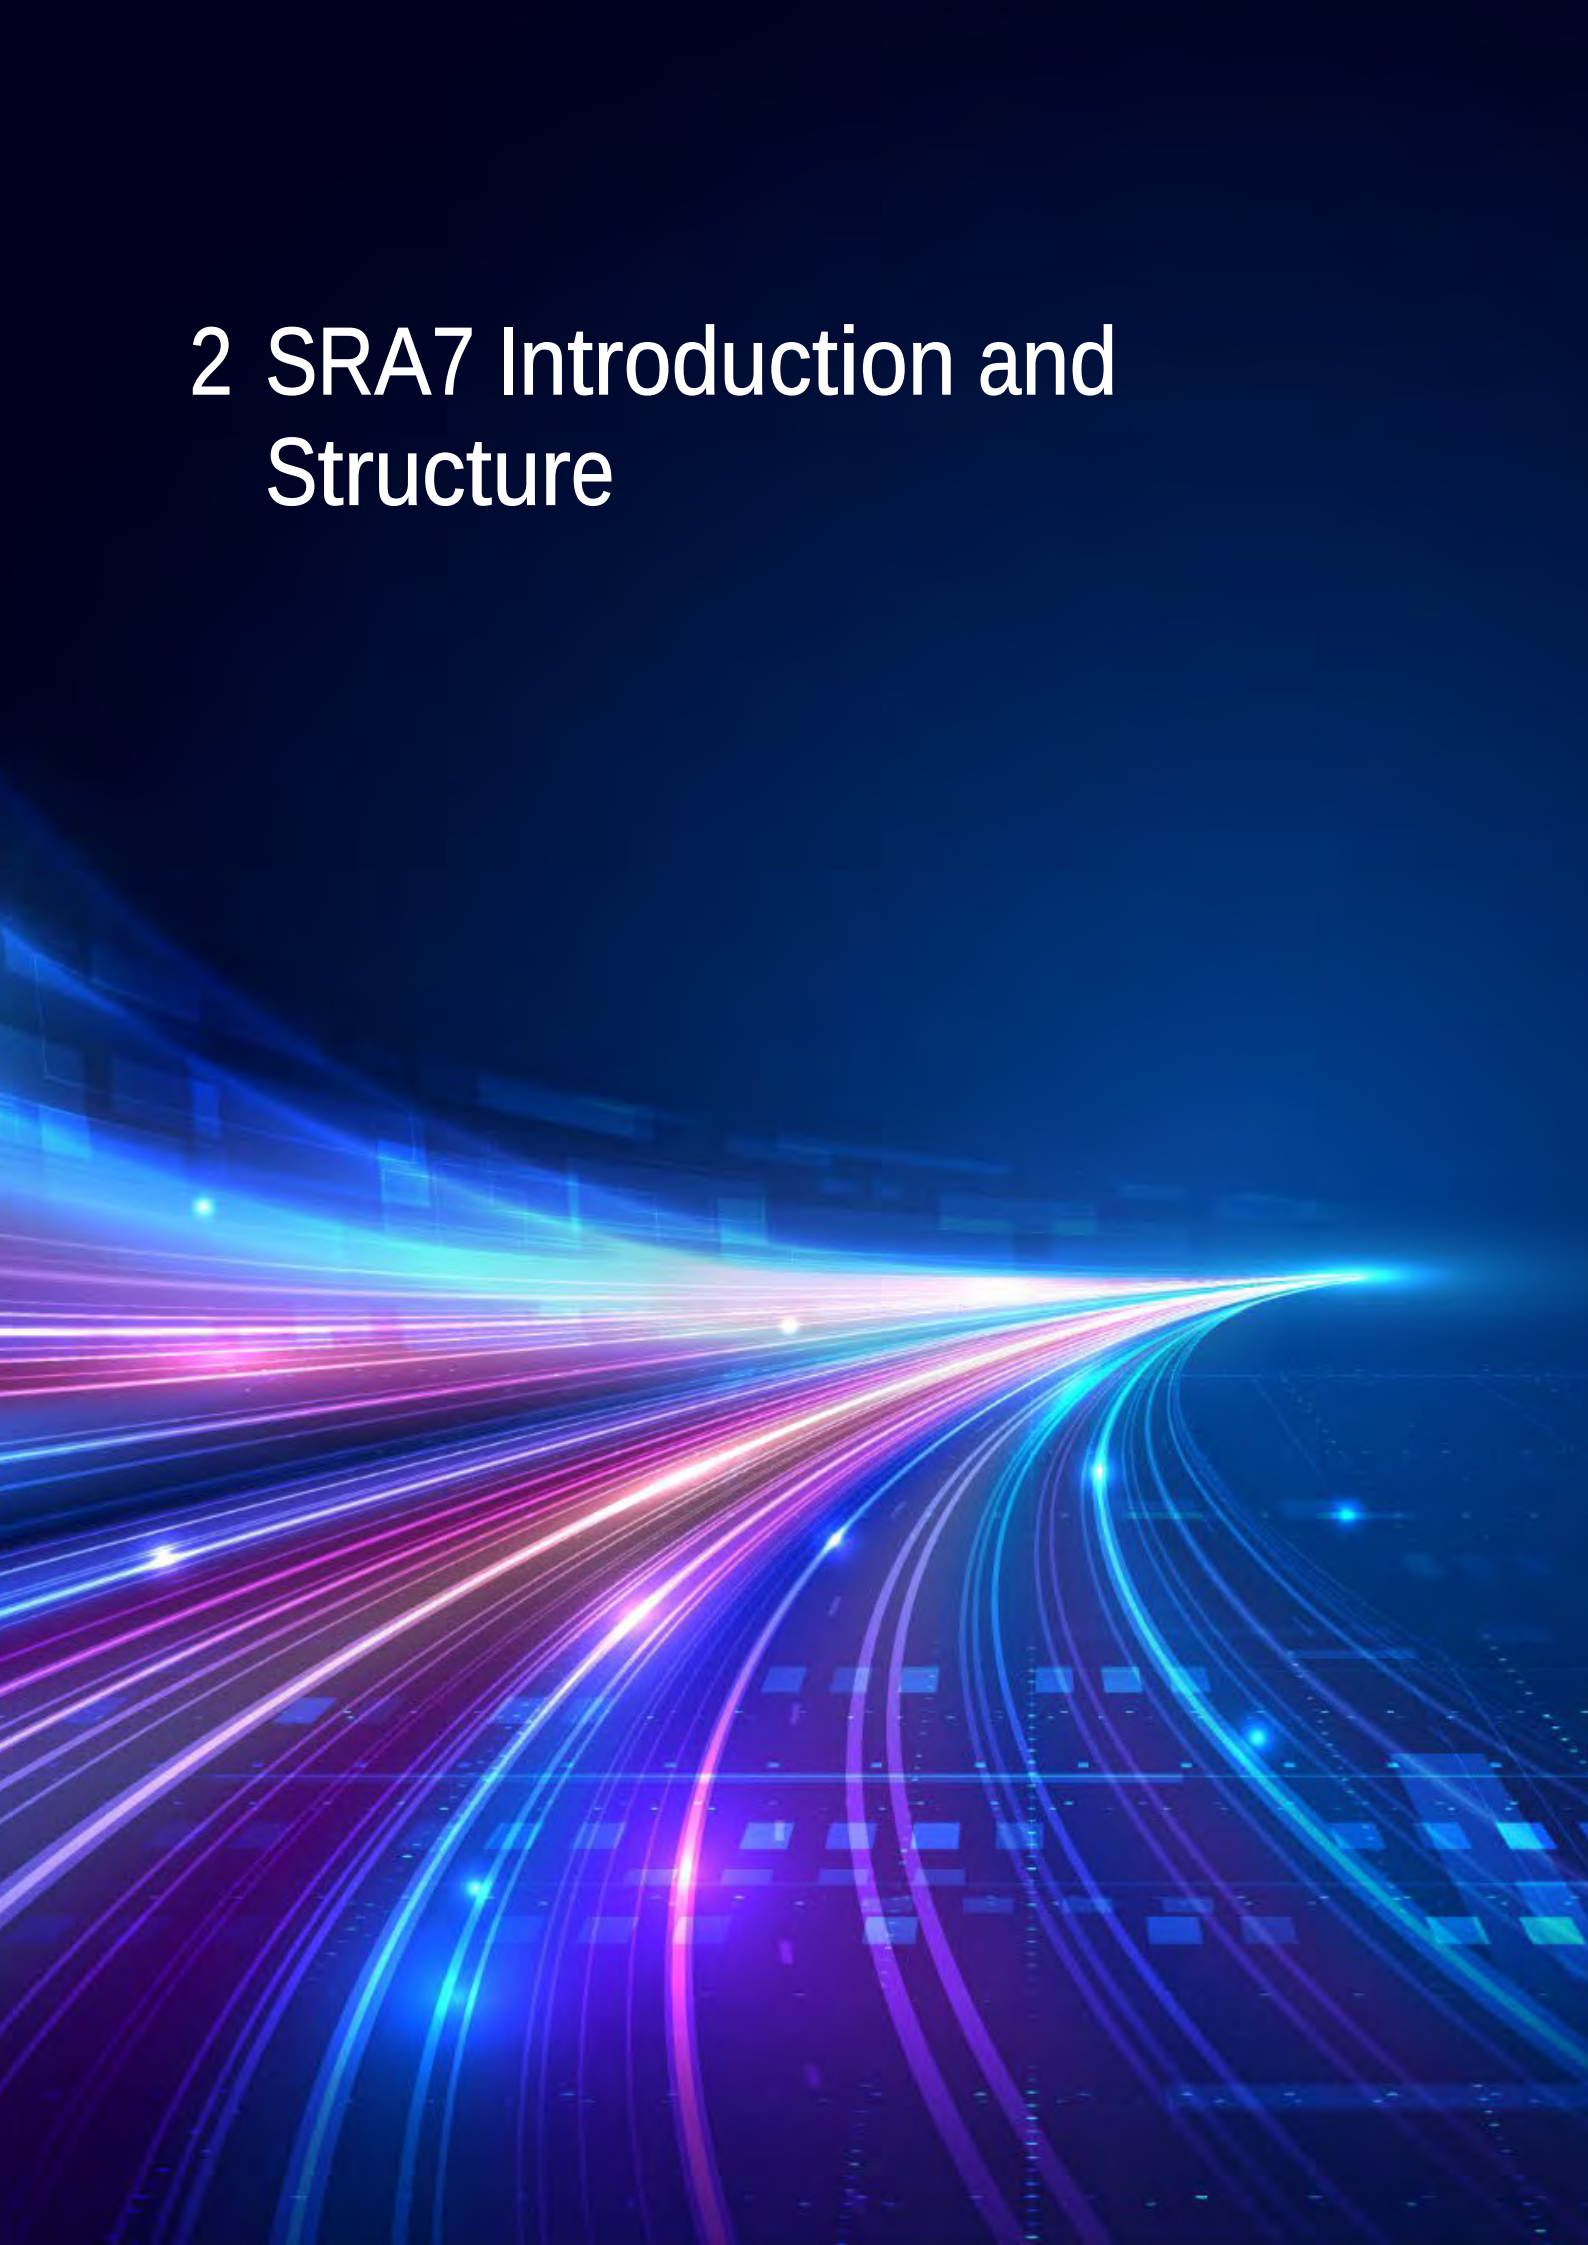
We show below the mapping of strategic recommendations and supporting white papers and our recommended estimated target year for concrete actions in that area.

No.	Strategic Recommendation	Supporting SRA7 White Papers	Indicative Target Year
Theme 1 – European Sovereignty, Industrial Adoption and Ecosystem Leadership			
R1	Strengthen European AI Sovereignty and Accelerate AI-Driven Scientific Discovery	HPC/AI Applications and Workflows; Mathematics and Algorithms; Industrial Usage	2029
R2	Make Post-Exascale a coordinated European mission	Non-Conventional Architectures and Technologies; Hardware Components; European Computing Continuum	2028
R3	Give European SMEs and industry easier access to Advanced Computing	HPC Technology SMEs in Europe; Industrial Usage; European Computing Continuum	2030
R4	Coordinate skills, governance and long-term ecosystem stewardship	Industrial Usage; HPC Technology SMEs in Europe; HPC–Quantum Computing Integration	2031
Theme 2 – Federated, Sustainable and Trustworthy Computing Infrastructure			
R5	Build a genuinely federated European computing continuum	European Computing Continuum; System Software and Management; Data Management, I/O and Storage	2030
R6	Establish Enforceable Interoperability Requirements in European HPC Procurement	European Computing Continuum; Programming Environments; Hardware Components	2028
R7	Put energy efficiency and sustainability into every investment decision	Energy Efficiency and Sustainability; System Software and Management; Mathematics and Algorithms	2028
R8	Make software provenance tracing and cyber-resilience mandatory	System Software and Management; Data Management, I/O and Storage; European Computing Continuum	2029
Theme 3 – European Computing Technologies and Architectures			
R9	Continue investments in European processors, accelerators and interconnects	Hardware Components; Programming Environments; HPC Technology SMEs in Europe	2030
R10	Establish a European chiplet and advanced-packaging ecosystem	Hardware Components; Non-Conventional Architectures and Technologies; HPC Technology SMEs in Europe	2031

No.	Strategic Recommendation	Supporting SRA7 White Papers	Indicative Target Year
R11	Create a modular experimental facility for disruptive computing	Non-Conventional Architectures and Technologies; Hardware Components; Mathematics and Algorithms	2029
R12	Advance Open, Secure and Scalable Quantum–HPC Integration	HPC–Quantum Computing Integration; Programming Environments; Non-Conventional Architectures and Technologies	2030
Theme 4 – Software, Algorithms and Intelligent Computing			
R13	Optimise data sharing and access, minimising data movement	Data Management, I/O and Storage; European Computing Continuum; Programming Environments	2029
R14	Fund sovereign European software factories	System Software and Management; Programming Environments; HPC Technology SMEs in Europe	2029
R15	Modernise resource management for dynamic workflows	System Software and Management; HPC/AI Applications and Workflows; Programming Environments	2029
R16	Advance Scalable, Energy-Efficient and Performance-Portable Software and Algorithms for Post-Exascale Computing	Mathematics and Algorithms; Programming Environments; System Software and Management	2030

Note: Target years indicate significant implementation milestones, not completion of ongoing research or policy activities. These are indicative planning estimates.

2 SRA7 Introduction and Structure



The European Strategic Research Agenda 7, or the SRA7, reflects the state of the European Advanced Computing (AC) ecosystem, which includes High Performance Computing and associated areas such as Artificial Intelligence and Quantum Computing - edited and published by the European Technology Platform for High-Performance Computing (ETP4HPC). The publication of Strategic Research Agendas has been one of the primary goals of the ETP4HPC, which is the largest industry-led think tank for High-Performance Computing in Europe, since its inception in 2012. The Strategic Research Agendas have been published every two years, with the last SRA6 published towards the end of 2024. As with the previous SRAs, a lot has changed over the last couple of years, and that primarily also reflects the timing of the publication of these agendas, which captures the state of European Advanced Computing technologies at two-year "checkpoints", which we think is a reasonable time window to see measurable changes in the HPC landscape.

2.1 So why does the SRA exist and what are the primary objectives in SRA7?

This European HPC Strategic Research Agenda is intended to be **the** go-to document to capture the state of, and provide a roadmap for Advanced Computing Technology Innovation in Europe. The SRA7 is also intended to be fully vendor neutral and reflects the joint contributions of top Advanced Computing experts in Europe coming from both large European HPC, AI and Quantum technology providers and SMEs, European research organisations, and European academic institutions. The SRA7 helps to provide a glimpse of the state of play in a single set of documents - a main report together with a series of white papers - to the entire technology ecosystem both within Europe and globally.

The SRA7 is also intended to guide and advise policymakers in Europe on the research and innovation priorities in Advanced Computing and to offer bold and actionable recommendations to help act as a guide.

2.1.1 SRA7, a community perspective

The SRA7 is an agenda that presents a holistic view of the Advanced Computing ecosystem in Europe. The agenda also provides global visibility to "Advanced Computing in Europe" corresponding to the various Research and Innovation topics and the associated ecosystem.

Advanced computing in general and HPC in particular consists of multiple research areas, many of which are evolving very fast. It is very hard to infer the developments in all these research areas in a single document, let alone keep pace with them all - whether programming models, hardware technologies, system software, AI, Quantum, HPC data management, system architectures etc. The SRA7 aims to connect all these domains into a coherent picture in its bundled set of documents. The SRA7 also provides a technology generalist view into the latest trends and themes in Advanced Computing. The work aims to help the technology community understand and structure the Advanced Computing world.

HPC/AI and Quantum use cases are evolving to keep pace and exploit the various new innovations in HPC software and hardware infrastructures. There needs to be effective coordination between the application communities and the technology developments and have a closed-loop feedback mechanism between the two so as to make sure that the applications get the very best from what the technologies have to offer. The SRA7 is a means of connecting the Advanced Computing applications and use cases and having them work closely with the technology experts.

Finally, HPC/AC is just one piece of the "Digital Continuum" which consists of various actors in the ecosystem such as those from the Edge, Cloud, Cybersecurity, AI/Robotics, wireless networking etc.

There is a need not just to keep track of the developments in these areas but also to closely synchronise and work with them to come up with research priorities in AI/HPC. For one, we don't want to reinvent in HPC the areas that are already well covered by these ecosystems. Dialogue and synchronisation with these communities help to develop coordinated efforts and action plans within Europe beneficial to society, industry and science. This is also one of the objectives of the SRA7 with a view to having an impact on the community and society.

2.1.2 SRA7, a policy makers' perspective

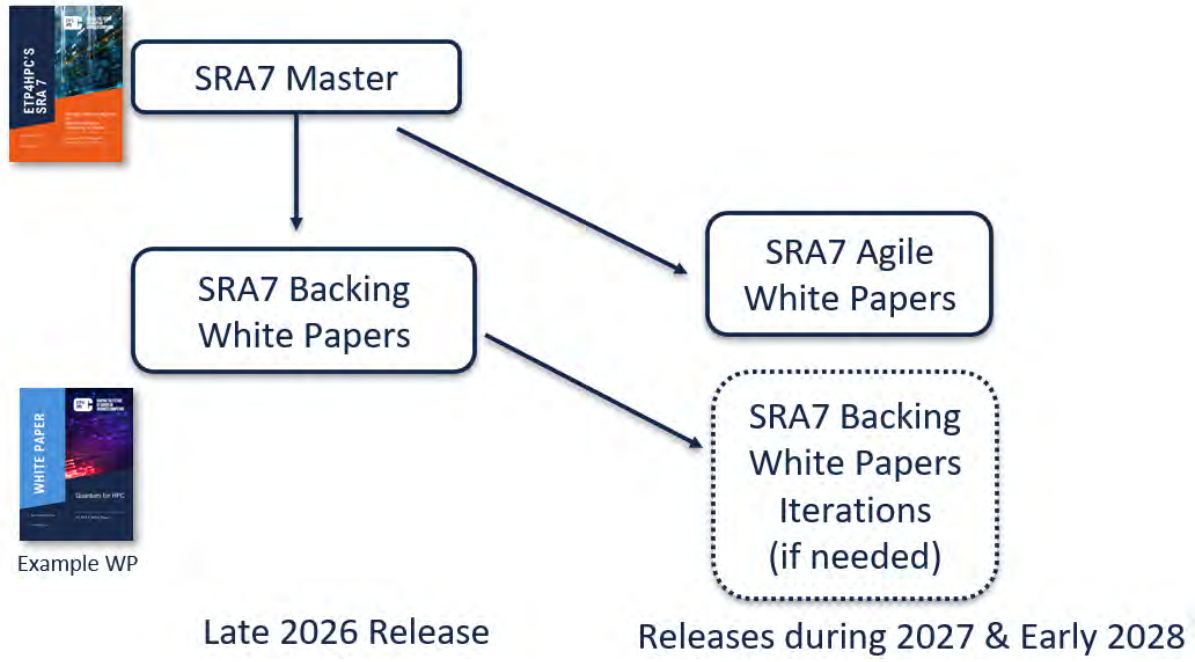
The SRAs in the past have influenced HPC Research and Innovation Work programmes under the Horizon 2020 Programme in the context of the cPPPs (contractual Public Private Partnerships). With the formation of the EuroHPC Joint Undertaking in 2018 and its RIAG (Research and Innovation Advisory Group), the SRAs have pivoted to accommodate this positive change and renewed focus in the European HPC ecosystem.

One of the primary goals of SRA7 is to provide technical advice and inputs into EuroHPC on Research and Innovation programmes and priorities. This will be done through the EuroHPC Research and Innovation Advisory Group (RIAG) with whom we have synchronised closely even while developing the SRA and we hope that this will help the EuroHPC Multi-Annual Strategic Programme (MASP). The SRA7 also aims to provide technical inputs and advice to the European Commission/DG-CNNECT (Directorate-General for Communications Networks, Content and Technology) on Advanced Computing Research and Innovation status and priorities. The SRA7 is also expected to be referred to by other Joint Undertakings in Europe in the technology domain to develop their respective Strategic Research Agendas.

2.1.3 SRA7 structure and release

We improved the structure and organisation of the SRA in the previous SRA6, where we broke down the entire work into detailed White Papers that go into the various subject areas, and a Master document that crystallises and summarises the key points and is fronted by bold actionable recommendations. This helps to preserve the various levels of focus needed by the readers – be it subject matter experts (who would go into the individual white papers in detail, for example), and technology generalists and policymakers (who might look into the Master) and more generally politicians and related administration experts who perhaps might just refer to the actionable recommendations at the very top of the Master. We have received very positive feedback from the community about the change in this approach from SRA6 onwards. We therefore preserve the same structure and organisation in SRA7.

The backing white papers associated with this Master are released almost at the same time or very close to the release of the Master. Then there are "agile" white papers that will be released during 2027 and early 2028 that cover developments in the ecosystem during that period. We reiterate this structure as below:

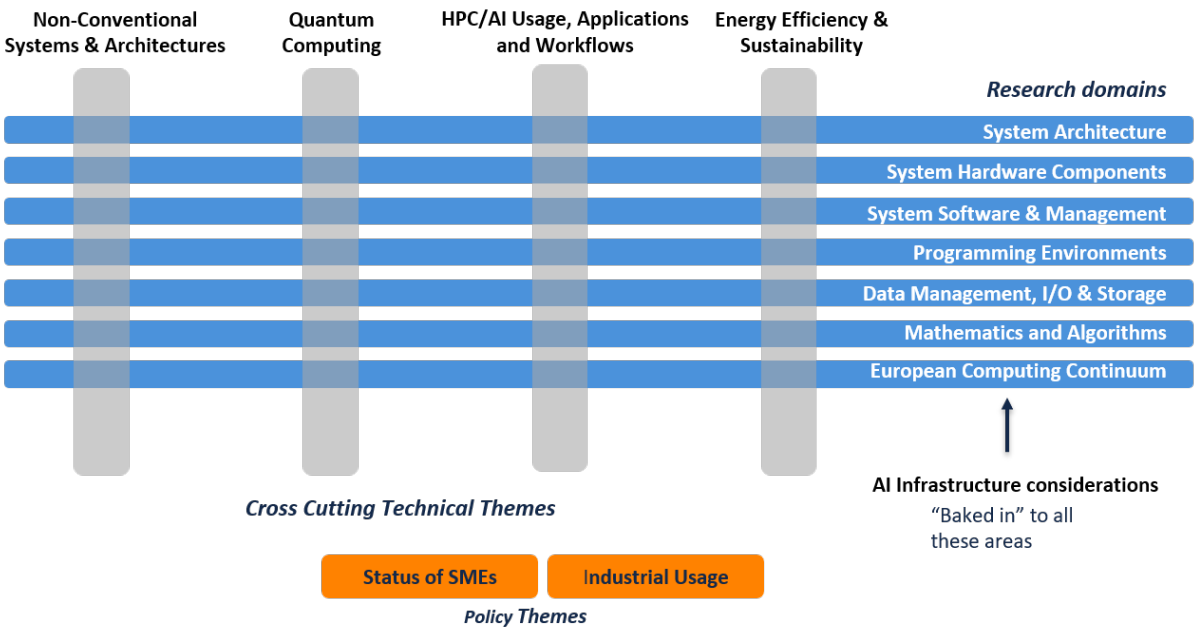


The SRA7 Backing White Papers will also be edited and reviewed as needed based on developments in the ecosystem during 2027 and 2028. We then head into the new SRA8 towards the end of 2028.

This flexible release mechanism, we believe, will keep the SRA current with the right level of detail as needed by different types of readers.

2.1.4 SRA7 white papers and master

The SRA7 white papers will cover 7 Research Domains, 4 cross-cutting technical topics and 2 Policy Themes (a total of 13 White Papers). We had AI and foundation models as a separate white paper in SRA6, but AI considerations will be addressed across all these topics and there will be no separate White Paper.



This document, the SRA7 Master, in addition to top-level recommendations, will cover the overall state of play of the European Advanced Computing ecosystem, will summarise the recommendations from each of the technical white papers and also include dedicated sections for the Status of SMEs and Industrial Usage.

2.2 Continued strong linkages with EuroHPC and the EC

We have continued to maintain strong links with the EuroHPC JU through our RIAG (Research and Innovation Advisory Group) members, the INFRAG, and the newly formed QTAG (Quantum Technology Advisory Group) who are also active contributors to the SRA. Our SRAs continue to provide foundational material for the RIAG's MASP (Multiannual strategic Plan) guiding the EuroHPC JU's actions and priorities. We have also discussed and provided specific inputs to the European Commission based on their requests for further information and SRA7 is expected to continue that activity, including the potential to spawn specific agile white papers that are needed to address, for example, policy priorities.

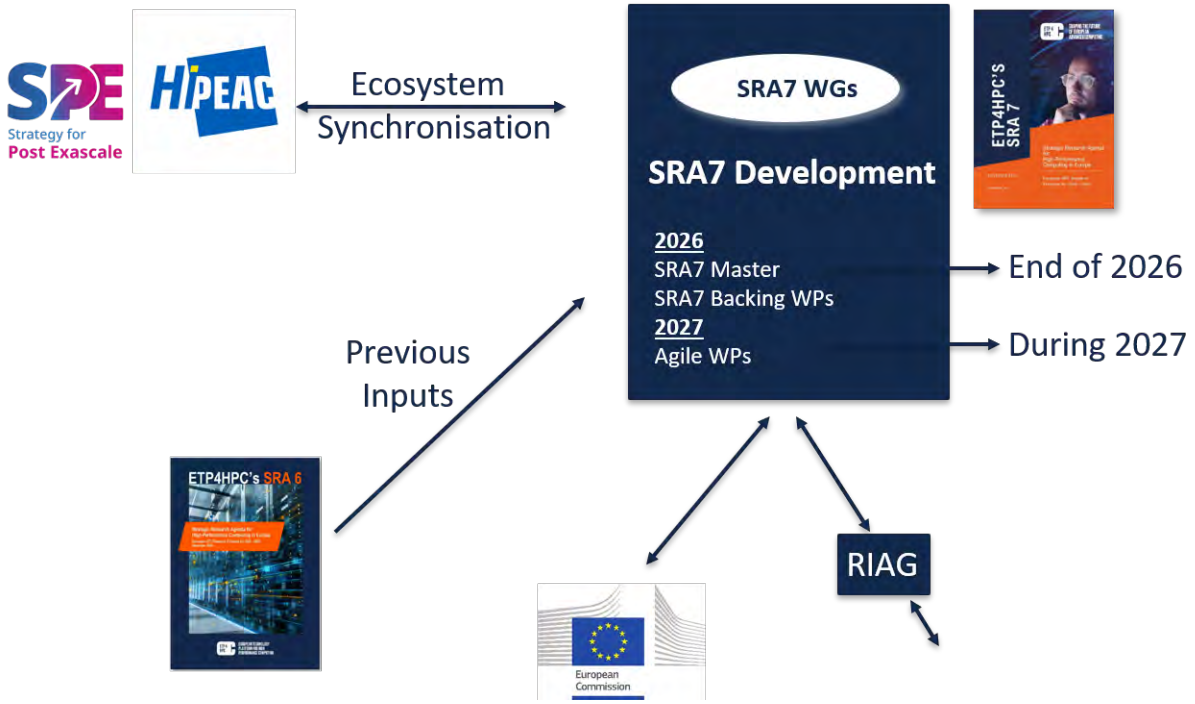
The SRA7 is prepared by independent experts from the ETP4HPC community; it does not necessarily reflect the positions of individual members of the association.

2.3 Continued strong engagement with the ecosystem communities

The SRA5 covered some of the early work done with the ecosystems adjoining HPC through the "Transcontinuum Initiative" (TCI). [7] The SRA6 extended and expanded this collaboration through the SRA6 teams' engagement with new initiatives such as the Computing Continuum Initiative (CCI) spearheaded by HiPEAC.

[7] https://www.etp4hpc.eu/2020/12/11/transcontinuum-initiative-joint-vision-document-signed-by-8-european-associations

We are taking ecosystem collaboration to the next level in SRA7. ETP4HPC is now on the board of HiPEAC and is driving the post-exascale European Roadmap as part of the EC-funded SPE-EU project. The outcomes of the SRA7 will feed into the post-exascale roadmaps developed by SPE-EU (Strategy Post-Exascale – EU) as well as the upcoming HiPEAC Vision 2027 developed by HiPEAC. HiPEAC continues to consult ETP4HPC SRA7 core members, and the SRA7 contributors also play a major part in the SPE-EU project. This indeed bolsters overall collaboration with hundreds of experts throughout Europe and helps to coordinate efforts from the other related communities to make Europe a very strong player in the global technology ecosystem.



2.3.1 Strong focus on community feedback

We will continue to gather feedback on the readership and obtain feedback on the SRA through various events and workshops ETP4HPC participates in and ETP4HPC webinars presenting the White Papers. The newly updated ETP4HPC website will track downloads of the SRA7 documents.

3 European Advanced Computing - Overall State of Play



Since the preparation of SRA6, the European high-performance computing ecosystem has entered a new phase, including the rapid global developments in AI (and agentic AI) driven science and innovation (e.g., Genesis Mission | Computational Science | NLR[8]). The policy objective is no longer limited to procuring world-class supercomputers and supporting their scientific use. EuroHPC is developing into a broader strategic computing infrastructure player that brings together high-performance computing, AI, quantum computing, data resources, networking, software, applications and skills. This reflects the growing importance of Advanced Computing to European competitiveness, technological sovereignty, scientific leadership and industrial innovation.

The period after November 2024, since we launched the last SRA, has been particularly significant. Europe has entered the exascale era with JUPITER; EuroHPC has begun building a distributed network of AI Factories and preparing much larger AI Gigafactories; several European quantum computers have become operational; and the EuroHPC Federation Platform has begun simplifying access to distributed European resources. At the same time, investments in European processors, accelerators, interconnects, system software, applications and skills have continued. The following developments provide the policy and technological context in which SRA7 has been prepared (after November 2024).

3.1 AI Factories and the rapid expansion of AI infrastructure

The most substantial change has been the expansion of EuroHPC's responsibilities in AI. In December 2024, EuroHPC selected the first seven AI Factories[9], located in Finland, Germany, Greece, Italy, Luxembourg, Spain and Sweden. A further six AI Factories were selected in March 2025[10], followed by six additional sites in October 2025[11]. This brought the network to 19 AI Factories. In parallel, 13 AI Factory Antennas were selected[12] to broaden geographical coverage and connect national AI ecosystems to the main factories.

Europe's ambition has subsequently moved beyond AI Factories towards substantially larger AI Gigafactories. The Commission's InvestAI initiative[13] aims to mobilise €200 billion for European AI, including investment in very large AI infrastructure. The AI Continent Action Plan[14] placed AI Factories, Gigafactories, data and skills at the centre of European AI policy. EuroHPC's mandate was consequently expanded through Council Regulation (EU) 2026/150[15], and the AI Gigafactories call was launched in July 2026[16]. These developments make AI a permanent and increasingly central pillar of EuroHPC.

[8] https://www.nlr.gov/computational-science/genesis-mission
[9] https://www.eurohpc-ju.europa.eu/selection-first-seven-ai-factories-drive-europes-leadership-ai-2024-12-10_en
[10] https://www.eurohpc-ju.europa.eu/eurohpc-ju-selects-additional-ai-factories-strengthen-europes-ai-leadership-2025-03-12_en
[11] https://www.eurohpc-ju.europa.eu/eurohpc-ju-selects-six-additional-ai-factories-expand-europes-ai-capabilities-2025-10-10_en
[12] https://www.eurohpc-ju.europa.eu/eurohpc-ju-selects-ai-factory-antennas-broaden-ai-factories-initiative-2025-10-13_en
[13] https://luxembourg.representation.ec.europa.eu/actualites-et-evenements/actualites/eu-launches-investai-initiative-mobilise-eu200-billion-investment-artificial-intelligence-2025-02-11_en
[14] https://commission.europa.eu/topics/eu-competitiveness/ai-continent_en
[15] https://www.eurohpc-ju.europa.eu/eurohpc-jus-mandate-expanded-under-new-regulation-amendment-2026-01-20_en
[16] https://www.eurohpc-ju.europa.eu/eurohpc-joint-undertaking-launches-ai-gigafactories-call-2026-07-30_en

3.2 Europe enters the exascale era

Europe has now crossed the threshold into exascale computing. JUPITER was inaugurated in September 2025[17]at Forschungszentrum Jülich. In November 2025, it became the first European system to break the exascale barrier in the November 2025 TOP500 ranking[18]. Together with LUMI, MareNostrum5 and Leonardo, JUPITER placed three EuroHPC systems among the world's ten most powerful supercomputers. In the same list, European systems occupied eight of the top ten places signifying its focus on energy efficiency and sustainability. The energy efficiency of JUPITER's architecture has also been demonstrated by the leading position achieved by its JEDI module in the Green500 list in June 2025[19].

The EuroHPC infrastructure has continued to expand beyond JUPITER. A procurement contract was signed for Alice Recoque, Europe's second exascale supercomputer[20], which will be hosted by GENCI and operated at CEA in France. New mid-range systems are also being deployed, including DAEDALUS in Greece[21], LEVENTE in Hungary[22], CASPIr in Ireland[23] and Arrhenius in Sweden[24].

A further notable development is the extension of the infrastructure towards industrial use. The INNOVATE consortium was selected[25] to host EuroHPC's first industrial-grade supercomputer in Bologna. This represents an important broadening of EuroHPC's mission from primarily scientific infrastructure towards systems and services designed explicitly for industrial innovation.

3.3 Quantum computing becomes an operational pillar

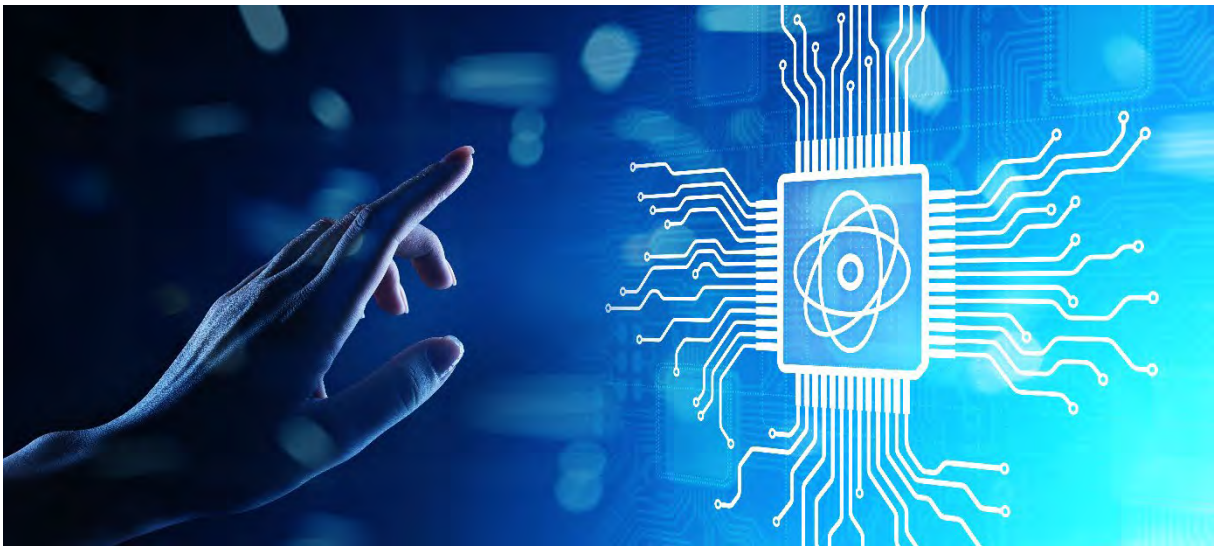
EuroHPC's quantum programme has progressed from hosting agreements and procurements to an operational European infrastructure. The programme deliberately supports multiple technological approaches, including superconducting circuits, trapped ions, neutral atoms, photonics and quantum annealing, and soon spin qubits. These systems are intended to be integrated with EuroHPC supercomputers, enabling hybrid quantum–classical workflows and allowing users to explore which technologies are best suited to different scientific and industrial problems.

[17] https://www.eurohpc-ju.europa.eu/jupiter-launching-europes-exascale-era-2025-09-05_en
[18] https://www.eurohpc-ju.europa.eu/jupiter-officially-propels-europe-exascale-era-2025-11-17_en
[19] https://top500.org/lists/green500/list/2025/06/
[20] https://www.eurohpc-ju.europa.eu/contract-signed-alice-recoque-europes-new-exascale-supercomputer-2025-11-18_en
[21] https://www.eurohpc-ju.europa.eu/eurohpc-ju-signs-procurement-contract-daedalus-supercomputer-2025-03-28_en
[22] https://www.eurohpc-ju.europa.eu/way-open-building-eurohpc-world-class-supercomputer-hungary-2025-07-09_en
[23] https://www.eurohpc-ju.europa.eu/way-open-building-eurohpc-world-class-supercomputer-ireland-2025-10-13_en
[24] https://www.eurohpc-ju.europa.eu/eurohpc-ju-inaugurates-arrhenius-new-mid-range-supercomputer-together-naiss-sweden-2026-09-08_en
[25] https://www.eurohpc-ju.europa.eu/innovate-consortium-chosen-host-eurohpcs-first-industrial-supercomputer-2024-12-04_en

PIAST-Q in Poland became the first inaugurated EuroHPC quantum computer[26], followed by VLQ in Czechia[27], Euro-Q-Exa in Germany[28], Lucy in France[29], EuroQCS-Spain[30] and SOL in Italy[31]. The Jade and Ruby processors developed through HPCQS have also been integrated into a federated hybrid HPC–quantum infrastructure[32]. SURF in the Netherlands and MeluXina in Luxembourg will also host spin-based quantum computers procured by EuroHPC.

A decisive transition occurred in June 2026 when EuroHPC opened access to its quantum computers[33]. Quantum computing has also been incorporated formally into EuroHPC's amended mandate, while the Quantum Excellence Centres QEX[34] and QEC4QEA[35] are supporting the development of applications and user communities. These activities complement the wider EU Quantum Europe Strategy[36], which seeks to connect European research leadership with infrastructure, industrialisation and skills.

At an administrative level, the Quantum Technology Advisory group (QTAG)[37] now forms a part of the EuroHPC JU RIAG. In April 2026, an independent and platform-agnostic advisory body established by the European Union was also launched: Q-CAB , the Quantum Computing Advisory Board.



[26] https://www.eurohpc-ju.europa.eu/inauguration-piast-q-leap-european-quantum-computing-2025-06-23_en
[27] https://www.eurohpc-ju.europa.eu/vlq-europes-next-leap-towards-quantum-powered-research-and-innovation-2025-09-23_en
[28] https://www.eurohpc-ju.europa.eu/inauguration-euro-q-exa-expanding-european-quantum-computing-infrastructure-2026-02-12_en
[29] https://www.eurohpc-ju.europa.eu/eurohpc-ju-inaugurates-quantum-computer-france-strengthening-europes-sovereignty-2026-04-14_en
[30] https://www.eurohpc-ju.europa.eu/eurohpc-ju-inaugurates-euroqcs-spain-barcelona-2026-05-28_en
[31] https://www.eurohpc-ju.europa.eu/inauguration-sol-new-quantum-computer-together-lisa-upgrade-leonardo-supercomputer-2026-06-11_en
[32] https://www.eurohpc-ju.europa.eu/inauguration-jade-and-ruby-quantum-processors-milestone-towards-hybrid-high-performance-computing-2025-11-13_en
[33] https://www.eurohpc-ju.europa.eu/eurohpc-ju-opens-access-its-quantum-computers-2026-06-25_en
[34] https://www.eurohpc-ju.europa.eu/qex-european-quantum-excellence-centre-drive-innovation-and-applications-2025-11-12_en
[35] https://www.eurohpc-ju.europa.eu/qec4qea-building-europes-one-stop-shop-quantum-enhanced-applications-2026-02-24_en
[36] https://www.europarl.europa.eu/legislative-train/theme-a-europe-fit-for-the-digital-age/file-european-quantum-act
[37] https://www.eurohpc-ju.europa.eu/about/governance/quantum-technologies-advisory-group-qtag_en

3.4 European technological sovereignty and the RISC-V ecosystem

Reducing strategic dependencies in processors, accelerators, networking and software remains a major European objective. EPI[38], the European Processor Initiative, supported, until 2026, a first level of ARM-based and RISC-V-based HPC European processor design. The DARE project[39] represents a significant further step in developing a European HPC and AI ecosystem based on the open RISC-V instruction-set architecture. Its scope extends beyond individual processor designs to accelerators, system architecture and the supporting software environment required for future exascale and post-exascale systems.

European sovereignty also depends on technologies for moving data efficiently between increasingly heterogeneous processors, accelerators, memory and storage. The NET4EXA project[40] is developing the next generation of European high-bandwidth, low-latency interconnect technology for HPC and AI systems. Together, DARE, NET4EXA and continuing European software-stack initiatives demonstrate a move towards a more complete European technology supply chain rather than reliance on isolated European components. The EUPEX project[41] which is now close to completion has started to develop an ecosystem (including the software stack and application co-design) around the European Processor ARM-based track. In parallel, the European PILOT project[42] aims to deliver the first All-European open source and open standard based software and hardware integrated HPC system by creating an autonomous set of accelerators designed, implemented, manufactured, and deployed in Europe.



[38] https://www.european-processor-initiative.eu/
[39] https://www.eurohpc-ju.europa.eu/advancing-european-sovereignty-hpc-risc-v-2025-03-06_en
[40] https://www.eurohpc-ju.europa.eu/net4exa-advancing-european-interconnect-hpc-and-ai-2024-12-13_en
[41] https://eupex.eu/
[42] https://eupilot.eu/

3.5 Federation of computing, data and software resources

The increasing diversity of EuroHPC systems makes federation essential. In December 2024, EuroHPC signed the contract to develop the EuroHPC Federation Platform[43]. A complementary contract was subsequently awarded to GÉANT to hyperconnect EuroHPC systems[44] through secure, ultra-high-bandwidth pan-European connectivity.

The first release of the Federation Platform[45] became available in April 2026. It provides a common access point and unified authentication, while progressively harmonising allocation, data-transfer, software-provisioning and user-support processes. Future releases are expected to incorporate AI Factories, quantum computers, data lakes and European data spaces. The Federation Platform is therefore an important transition from a collection of nationally hosted machines towards a more coherent European computing continuum.

3.6 Sustainability, software and applications

Energy efficiency has become a system-wide requirement rather than simply a hardware optimisation objective. It must be considered across architecture, workload scheduling, resource management, data movement, software and data-centre operation. The EuroHPC-funded SEANERGYS project[46] is addressing this challenge by developing software for more energy-aware operation of supercomputers and for coordinating energy use across applications, systems and data-centre resources.

Investment in applications and software remains equally important because infrastructure alone does not create scientific or industrial value. EuroHPC has launched a new call for Centres of Excellence and Lighthouse Codes[47], while projects such as MINERVA[48] connect HPC capabilities to AI communities. The European HPC Application Support Portal[49] and the HPC in Europe Portal[50] seek to make expertise, services and application support more accessible across Europe.

3.7 Skills, widening and international cooperation

The availability of people with advanced HPC, AI and quantum expertise remains a critical constraint. EuroHPC is continuing the EUMaster4HPC programme for 2026–2032[51], has introduced the HPCTRAIN

[43] https://www.eurohpc-ju.europa.eu/paving-way-eurohpc-federation-platform-2024-12-19_en
[44] https://www.eurohpc-ju.europa.eu/eurohpc-ju-awards-contract-geant-hyperconnect-european-supercomputers-2025-09-22_en
[45] https://www.eurohpc-ju.europa.eu/first-release-eurohpc-federation-platform-streamline-access-europes-supercomputing-resources-2026-04-15_en
[46] https://www.eurohpc-ju.europa.eu/new-eurohpc-project-seanergys-advances-energy-efficient-supercomputing-2025-06-25_en
[47] https://www.eurohpc-ju.europa.eu/new-eurohpc-call-strengthen-europes-leadership-hpc-applications-2025-06-10_en
[48] https://www.eurohpc-ju.europa.eu/advancing-ai-eurohpc-minerva-project-2025-02-13_en
[49] https://www.eurohpc-ju.europa.eu/european-hpc-application-support-portal-empowering-researchers-across-eurohpc-systems-2025-10-01_en
[50] https://www.eurohpc-ju.europa.eu/castiel-2-project-launches-hpc-europe-portal-europes-unified-gateway-supercomputing-growing-step-2025-10-31_en
[51] https://www.eurohpc-ju.europa.eu/eumaster4hpc-programme-continues-advancing-hpc-skills-europe-new-phase-2026-2032-2026-07-29_en

professional traineeship programme[52], and is funding a third phase of EuroCC and CASTIEL[53]. These initiatives provide complementary routes for university education, professional development, user support and wider adoption by industry and public administrations.

The geographical and international reach of the ecosystem has also expanded. Albania and Moldova joined EuroHPC, while Switzerland formally rejoined[54]. In October 2026, Canada joined the EuroHPC JU to become its 39th participating state[55]. These member states can also now benefit from the recommendations in this SRA. International projects include GANANA for EU–India cooperation[56], HANAMI for EU–Japan exascale collaboration[57], and Q-Neko for EU–Japan cooperation in hybrid quantum–HPC algorithms[58]. Such partnerships provide access to complementary knowledge while European technological capability and sovereignty remain central objectives.



[52] https://www.eurohpc-ju.europa.eu/european-hpc-professional-traineeship-programme-hpctrain-strengthen-europes-supercomputing-skills-2026-03-06_en
[53] https://www.eurohpc-ju.europa.eu/eurocc-3-and-castiel-3-further-supporting-existing-european-network-national-competence-centres-hpc-2026-08-31_en
[54] https://www.eurohpc-ju.europa.eu/switzerland-officially-rejoins-eurohpc-joint-undertaking-2025-11-11_en
[55] https://www.eurohpc-ju.europa.eu/canada-becomes-39th-participating-state-eurohpc-joint-undertaking-2026-10-30_en
[56] https://www.eurohpc-ju.europa.eu/new-eurohpc-project-strengthening-eu-india-ties-ganana-project-2025-02-14_en
[57] https://www.eurohpc-ju.europa.eu/riken-r-ccs-and-eurohpc-sign-landmark-letter-intent-empower-exascale-collaboration-hanami-project-2025-06-11_en
[58] https://www.eurohpc-ju.europa.eu/eurohpc-ju-launches-q-neko-project-strengthen-eu-japan-collaboration-quantum-computing-2026-07-06_en

3.8 Conclusion

The developments since SRA6 demonstrate a considerable broadening and maturation of the European HPC ecosystem. Europe now possesses operational exascale capability, is deploying a continent-wide network of AI Factories, is preparing AI Gigafactories (preparing for the era of Agentic AI), and has moved from experimental quantum procurements towards an accessible hybrid HPC–quantum infrastructure. The Federation Platform and high-capacity network investments are beginning to connect these resources, while ARM-based or RISC-V processors, European interconnects, software, applications and sustainability projects are strengthening the underlying technology base.

The principal challenge for the next period is therefore not simply to procure additional machines. It is to transform a diverse portfolio of infrastructures and projects into a coherent, sustainable and usable European computing ecosystem, preparing Europe to exploit Agentic AI and emerging post-exascale computing paradigms[59] (R1, R2). This requires interoperability and federation across HPC, AI, quantum and data resources (R5, R6, R12); portable, resilient and energy-efficient software and algorithms (R7, R14, R16); long-term support for applications and software ecosystems (R14, R15, R16); closer engagement with industry and SMEs (R3, R9–R11); and a sufficiently large, skilled and mobile workforce (R4). Together, these priorities provide the foundation for a coordinated European approach to post-exascale computing.

The future of scientific discovery will be shaped not merely by those who accumulate computing power, but by those who integrate it effectively.

[59] We refer to the recommendations from the "Executive Summary of Recommendations" in this paragraph



4 Advanced Computing: R&I Recommendations

This chapter summarises recommendations from the various research domains and cross-cutting technical themes as were defined in the Executive Summary of Recommendations. The policy themes are dealt with in separate chapters. More details backing up these recommendations can be found in the individual white papers which we encourage the reader to refer. **If not defined inline, the acronyms are defined in the glossary at the beginning of the report**.

Some overlaps exist between these recommendations, reflecting the perspectives and priorities of the independent working groups. These overlaps have been intentionally preserved to recognise the cross-cutting nature of the challenges and highlight areas where coordinated action across multiple research and technology domains is necessary.

4.1 Research Domain Recommendations

4.1.1 System Architecture

Quantum-Classical Integration Frameworks

Develop vendor-neutral interfaces for quantum device integration into HPC and AI workload managers, including calibration-aware scheduling, fidelity-aware resource allocation, and standardised telemetry. Support for QDMI/QRMI interoperability and hybrid quantum-classical workflow orchestration should be prioritised across EuroHPC facilities.

Co-design Methodologies for Post-Exascale Systems

Establish co-design frameworks that treat compute, memory, storage, networking, cooling, and power delivery as interdependent system components. Digital twin environments and AI-assisted simulation should support continuous architectural optimisation across the full infrastructure stack, including facility-level thermal and energy constraints.

Pan-European Energy Monitoring and Idle Compute Reduction

Define common data models, retention policies, and cross-site telemetry interfaces to enable federation-aware energy accounting. Coordinated action through EuroHPC should target unified observability across sites, enabling workload routing based on real-time energy conditions and eliminating idle compute waste at scale.

Open Standards for Heterogeneous Orchestration

Invest in open, interoperable APIs and programming models spanning CPUs, GPUs, AI accelerators, and quantum resources. Standardised scheduler interfaces, portable middleware, and, federated identity and policy frameworks are prerequisites for avoiding architectural lock-in and enabling workload portability across European HPC and AI ecosystems.

Software-Defined Infrastructure and AI-Native Operations

Advance declarative infrastructure management, GitOps-based provisioning, and AI-assisted anomaly detection and capacity planning. Autonomous operational capabilities - including predictive failure

analysis, energy-aware scheduling, and self-healing orchestration - should be treated as first-class infrastructure requirements rather than operational add-ons.

Federated Data Management and Sovereignty

Develop infrastructure for compute-to-data scheduling, distributed metadata catalogues, and federated storage namespaces that satisfy EU data residency and regulatory requirements. Data provenance, immutable audit logging, and policy-aware orchestration across jurisdictional boundaries must be co-designed with governance frameworks aligned to the EU AI Act and GDPR.

Usability and Democratisation

Invest in abstraction layers, conversational orchestration interfaces, and harmonised software environments that reduce the specialist expertise required to access federated HPC and AI infrastructure. Skill portability across sites, AI-assisted onboarding, and integrated training workflows are strategic requirements for expanding the European HPC and AI user base.

4.1.2 System Hardware Components

Further enhancing "Byte per Flop" ratio

Diverse heterogeneous memory solutions are available now (MRDIMM, HBM, DDR, etc). HBM is costly, difficult to buy currently and is power intensive. Alternatives and usage of such heterogeneous memories need to be further explored, Capabilities of technologies such as PIM has to be addressed to make PIM a viable proposition for HPC applications.

Coupling nodes into virtual "supernodes", a cluster composed of several nodes closely coupled together and virtually sharing resources (memory, data coherency) thanks to high bandwidth networks, guarantees high bandwidth and latency. Appropriate Interconnects for such "supernodes" needs to be looked into. We need to consider various evolutions of interconnects such as Infiniband, BXI, NVLink, Omni-path and the newly published open standards Ultra Ethernet (UEC 1.0, June 2025) and UALink (200G 1.0, April 2025; 2.0, April 2026) evolving into Ultra Ethernet (UEC). As AI models grow in complexity, the number of NICs per node is increasing for creating such "supernodes", further emphasizing the need for efficient data movement, memory coherence and orchestration. In this context, existing concepts implemented by technology vendors and associated standards continue to be very promising in the realm of Unified Memory architectures and needs to be further explored. Transparent management of memory and data movements in these supernodes with specialized network controllers (Infrastructure Processing Units, Data Processing Units, etc.) and computation on the data "in-transit" needs to be further explored. Software to support all hardware innovations in as transparent as possible fashion, leading to real hardware abstraction for the non-system programmer needs to be considered.

RISC-V

Continue European funding programs addressing RISC-V challenges for HPC applications. This should include the support of startups and innovative SMEs in this field, by easing the pain of very costly leading-edge design and silicon required for HPC applications.

Accelerators

Research and innovation should place greater emphasis on specialised, inference-optimised accelerators, addressing the rapidly growing demand for energy-efficient, high-performance AI inference across HPC, AI Factories, and data centres. This includes advancing dedicated inference architectures, NPUs, FPGAs, IPUs, DPUs, and accelerators for specialised algorithms such as digital annealing, while leveraging existing European capabilities in accelerator design. Hardware support for low-precision AI formats (e.g., MXFP4, MXFP6, MXFP8, and NVFP4) is becoming increasingly mature, but significant challenges remain in software frameworks, compiler support, and efficient hardware utilisation.

Variable Precision computing

Considering the AI driven trend towards lower precision FP support provisions in hardware, there is a need to look at how existing applications can move towards using this. How HPC adapts to lower precision AI "ecosystem" needs to be addressed. On the side of precision, higher resolution computation units (dealing with floating point representation using a large number of bits - e.g. mantissa coded up to 512 bits), could reduce the number of iterations of algorithms, even reaching convergence which is not possible with FP64. There is also a need for close numerical analysis and tools to move this area of variable precision (lower and higher than FP32) forward.

Integration of computing, memories and interconnects within the same package

A chiplet and interposer technology ecosystem should be created (interoperability between providers of chiplets, common HW interfaces, etc) with some efforts already in this direction. This is an action that should be promoted by Europe, allowing its diversity of SMEs, companies, start-ups to propose innovative IPs that could end-up in real HPC machines at minimum cost using this approach of interposers supporting a variety of IPs from different providers. The difficulty will be to convince non-European providers (High performance cores, memories) to provide their IPs in chiplet compatible form. The publication of UCIe 2.0 (August 2024) and UCIe 3.0 (August 2025), and the growth of the UCIe consortium to over 140 member companies suggests that the need for this is accelerating. Development of active interposers (power conversion, health monitoring, IOs) and photonic interposers to improve bandwidth per Watt is also needed. In the area of NOCs (Network on Chips), a 2D mesh is used in many existing NoCs. However, 3D integration should be considered for large future systems to go beyond 2D mesh capabilities.

Reliability

There is a need for robust real-time health monitoring systems (due to the increase of components in post-exascale systems) to detect corruptions and look into various methodologies (including those driven by AI) to handle them. The Open Compute Project's Silent Data Corruption initiative (AMD, Arm, Google, Intel, Meta, Microsoft, NVIDIA) is now the natural forum for this cross-industry work and should be tracked and, where possible, engaged with by European stakeholders.

Security

Continued exploration of DPUs & IPUs to offload a wide range of security tasks (is necessary, whilst noting that their interoperability with existing hardware/software remains a concern. Post-quantum cryptography standards such as NIST's (ML-KEM/FIPS 203, ML-DSA/FIPS 204, SLH-DSA/FIPS 205, August 2024) should be evaluated, given the multi-year lead time typically required for cryptographic migration.

Supply chain considerations

Europe should undertake a comprehensive assessment to strike the right balance between strengthening European supply-chain capabilities and diversifying sourcing through trusted international partners. In particular, European and national R&I programmes should prioritise investments in advanced packaging and chiplet integration technologies; support the development and industrial adoption of European memory technologies; promote access to EDA (Electronic Design Automation) tools through strategic partnerships, open alternatives and reduced vendor dependencies; and strengthen European semiconductor manufacturing capabilities through coordinated investment and procurement, building on the European Chips Act. These actions should be coordinated by the European Commission, Chips JU, Member States and industry, establishing a European chiplet and advanced-packaging ecosystem.



4.1.3 System Software and Management

Establish European "software factories" for continuous integration and deployment

Invest in shared, scalable CI/CD infrastructure — "software factories" — that build, integrate, test, and deploy HPC and AI/ML software stacks across the ecosystem, at both small/medium and large scale. Essential components (e.g., EasyBuild, Spack, GitLab, Podman, EESSI[60]) are mature and in wide use, yet assembling a robust, scalable CI/CD capability remains error-prone and maintenance-heavy over the long term.

The target should be to have integration and CI/CD tooling for large codebases with wide dependency graphs and complex configuration spaces demonstrated as prototypes in operational environments by

[60] https://www.eessi.io/

2028; European governance of at least one production-grade software factory serving EuroHPC systems should be made available.

Build verifiable provenance and supply-chain integrity into the software factories

Make code and data provenance a first-class, automated output of the build-and-deploy pipeline rather than a manual afterthought. Generate Software Bills-of-Materials (SBOM; SPDX / CycloneDX) and signed build attestations (in-toto, SLSA levels, Sigstore-style signing) for every artifact, anchored to hardware roots of trust (e.g., Caliptra / DC-SCM attested-boot → attested-build → attested-run chains). Extend this to converged workflows with W3C PROV / RO-Crate data lineage aligned to FAIR/EOSC and integrate continuous vulnerability (CVE) scanning into the pipelines. A sovereign ecosystem must be able to prove what its software is made of and where its data came from.

The Indicative target should be to have SBOM plus signed attestation for all software delivered through EuroHPC-hosted software factories by 2028; reproducible builds demonstrated for a reference stack (e.g., EESSI); and; provenance-capture overhead kept within an agreed budget.

Silent Data Corruption detection at scale

Treat correctness as a system-level property on a par with performance. Fund low-overhead, in-production SDC detection (in the spirit of lightweight, workload-co-located screening) and the embedding of Algorithm-Based Fault Tolerance (ABFT) into widely used scientific libraries and broaden regression testing in the software factories from performance-only to correctness and resilience regression. Detection must respect rigid HPC time budgets — avoiding blanket DMR/TMR redundancy — and must account for heterogeneity (CPU, GPU, FPGA, SmartNIC, AI accelerators) and for error propagation through MPI collectives.

The target should be to have reference implementations of low-overhead SDC screening integrated into at least two European scientific libraries, with an agreed overhead ceiling.

Enable trustworthy AI-in-the-loop orchestration, backed by open European benchmarks

Support the safe integration of AI — including LLM-based agents — into resource scheduling, software automation and cluster administration, with sandboxing, static analysis and human-in-the-loop approval before any autonomous action, and with on-premises or distilled inference for air-gapped and constrained sites. Critically, fund the creation of open, labelled European benchmark suites and datasets — for SLURM anomaly detection, natural-language-to-command translation, incident resolution and multi-objective scheduling — for which no widely adopted, openly licensed suites are yet available, and without which progress cannot be objectively measured or compared across systems. Provenance should extend to AI-generated code and configurations, so that autonomous actions remain auditable.

Indicative target should be to have at least one openly licensed benchmark suite and reference evaluation harness released by 2028, and a documented safety envelope for autonomous actions.

Advance energy and carbon efficiency through AI-assisted profiling and standard telemetry

Reduce the manual effort and specialist expertise currently needed to reach efficient execution by investing in AI-assisted performance and energy profiling and autotuning, at both application and system level, and in standardised energy telemetry and reporting across the stack. Efficiency gains must

be reproducible and comparable across sites and architectures, so that improvements can be verified rather than merely claimed.

Indicative target should be to have a standardised energy-telemetry interfaces adopted by European sites, and AI-assisted autotuning demonstrated to shorten time-to-efficient-configuration for representative codes.

Deliver standard-based federated services for access, identity, and deployment

Provide standards-based federated services for access to, and deployment of, applications across the compute continuum and multiple compute facilities. This includes federated identity and Single Sign-On (SSO), policy-driven and self-service resource provisioning through HPC gateways/portals, and governance and usage-control policies for the more flexible, volatile perimeters of federated data, including the means to audit — and, at defined policy-enforcement points, to constrain — the future usage of trained models.

The target should be to have interoperable federated-identity and provisioning interfaces piloted across at least two European centres by 2029.

Evolve resource and job management towards dynamic, live orchestration

Evolve Resource and Job Management Systems (RJMS) towards the dynamic, live orchestration that new, complex application workflows demand — a capability currently in tension with the static resource-management practices of today's supercomputers, and among the clearest open issues in the compute continuum. This includes malleable and evolving jobs, dynamic node pools, and container-native operation, building on standard interfaces such as PMIx, together with co-designed interaction with workflow engines so that resource allocations can follow workflows at run time. Dynamic orchestration is a prerequisite both for the AI-in-the-loop scheduling to act on its decisions and for the federated services to offer elastic access; it is nevertheless a distinct engineering problem in its own right, independent of whether AI is in the loop.

Indicative target should be to have RJMS extensions for dynamic orchestration (malleability, dynamic node pools) demonstrated with at least one widely used workflow engine on a EuroHPC system by 2029.



4.1.4 Programming Environments

The first five recommendations below address programming and execution capabilities; The rest provide shared evaluation, deployment and maintenance foundations. Research funding should be complemented by multi-year support for integration, testing services and critical software maintenance. Alongside technical results, projects should report use by independent application teams, released integrations and arrangements for continued support.

Strengthen portable compilers, abstractions and numerical libraries

Develop programming abstractions with compiler and runtime support, together with representative applications and numerical libraries. Priorities include vectorisation, accelerator offload, distributed dense and sparse computations, mixed-precision methods, and ways to communicate data layout and numerical requirements between software layers. Support modernisation of legacy C, C++ and Fortran applications alongside high-level languages and DSLs. Assess emerging platforms, including RISC-V, through complete application paths from compilation to deployment, and contribute improvements to maintained upstream projects.

An Indicative target would be to demonstrate maintained application suites on CPU and accelerator platforms from at least two processor or accelerator vendors, including hardware designed by at least one European-designed platform vendor and at least one emerging HPC architecture, with published coverage, performance relative to tuned baselines and porting effort.

Coordinate data placement, communication and I/O across workflows

Develop methods that jointly optimise where data resides, where computation executes and how transfers are scheduled across memory, storage and networks. Combine topology-aware placement, asynchronous communication and I/O, data reuse, in situ analysis and suitable compression with information about dependencies and access patterns. Evaluate complete simulation and analysis workflows, including coupled AI stages, and account for contention, capacity limits and numerical accuracy. Where workflows span edge, cloud and HPC resources, include transfer costs and restrictions on data location.

An Indicative target would be to demonstrate reduced data movement and improved time or energy to a validated result in at least three representative workflows on two heterogeneous systems, compared with documented baseline configurations.

Make dynamic execution composable and resilient

Build on Europe's established strengths in resource sharing, distributed task execution and malleability, and support the evolution of standards such as MPI to improve support for dynamic resource management. Develop and adopt common interfaces through which applications, runtimes and resource managers coordinate communication progress, resource sharing, changing allocations and recovery. Shared low-level runtime services integrated with task-based runtimes provide a basis for cooperation among programming environments. Prioritise interoperability between independently developed components and integration with operational resource managers as requirements for wider adoption. Combine malleability and workflow elasticity with policies for power limits, urgent workloads and failures, addressing consistent state and data redistribution. Evaluate application performance, fairness and system throughput using realistic workload mixes and deployment pilots.

An Indicative target would be to validate at least two independent runtime–scheduler combinations integrating components from different development teams, demonstrating allocation changes and controlled failure recovery in operational pilots, with published correctness evidence, recovery times, adaptation overheads and effects on shared-system throughput. Include use by application teams outside the component-development projects and document the released integrations and support arrangements needed for continued operation.

Establish interoperable workflows involving simulation, AI and quantum services

Develop reusable interfaces for coupling established HPC applications to training, inference, differentiable components and quantum execution. Prioritise data representation, memory ownership, device placement, synchronisation, error handling and resource coordination. Evaluate both migration to differentiable toolchains and integration of automatic differentiation into existing languages and solvers. For HPC–QC workflows, focus on demonstrable orchestration capabilities, resource use and integration costs; the programming-environment target should not depend on a prediction of quantum advantage.

An Indicative target would be to validate at least three simulation–AI workflows using established HPC codes, including one differentiable workflow, and demonstrate a hybrid workflow connected to a QPU service, with published integration effort and data-transfer or orchestration overheads.

Develop dependable AI assistance for HPC modernisation and optimisation

Support AI-assisted refactoring, porting, documentation, test generation and optimisation through controlled access to compilers, profilers, schedulers and representative hardware. Combine generated transformations with independent tests, numerical criteria, concurrency analysis and human review. Evaluate maintained, multi-file applications and follow-up maintenance tasks, including MPI+X and accelerator code, rather than relying solely on kernels. Record model and tool versions, search budgets and change provenance, and assess confidentiality, software security and licensing requirements. Measure gains after including review, validation and computational costs.

An Indicative target would be to independently evaluate assisted development on maintained MPI+X applications across at least two hardware platforms, reporting validated outcomes, developer and reviewer effort, introduced defects, performance and total search cost against documented human-led baselines.

Make execution observable and performance evidence comparable

Develop interoperable instrumentation and analysis that connect application phases and source code to tasks, communication, device activity, data movement and scheduler decisions. Couple these capabilities to systematic benchmarking frameworks, with versioned definitions, representative application suites, repeated measurements and machine-readable results. Agree protocols for performance portability, scaling, energy use and developer productivity, while preserving platform-specific reference measurements and exposing unsupported configurations. Measure diagnostic overhead and ensure that explanations remain useful at scale.

An Indicative target would be to establish a common evaluation protocol by the end of 2028 and, by 2031 to publish reproducible results across at least three hardware architectures, including performance-portability coverage, measurement uncertainty, diagnostic overhead and developer-effort data.

Enable federated validation and reproducible deployment

Develop supported interfaces for unattended builds, tests and benchmarks across HPC facilities, with appropriate authentication, accounting and scheduling. Integrate package and environment management, containers, dependency inventories, provenance and archived build configurations. Validate host-driver, communication-library and runtime compatibility, and provide authenticated artefacts and controlled access for automated tools. Coordinate these capabilities with facility operators so that workflows are usable under real security, queueing and resource-allocation policies.

An Indicative target would be to operate authorised unattended validation pipelines across at least three participating HPC facilities, with reconstructable environments, retained provenance and documented procedures for diagnosing and correcting cross-site failures.

Sustain critical open software and the expertise to maintain it

Treat maintenance, documentation, training and upstream integration as funded parts of the R&I programme, preserving European expertise and enabling successful research to progress towards wider operational adoption. Identify critical compilers, libraries, runtimes, workflow tools and development infrastructure; establish ownership, release policies, dependency support and security-update responsibilities. Encourage open interfaces and multiple implementations, with governance that allows users and maintainers to influence priorities. Support research software engineers, contributor development and migration assistance so that successful prototypes become maintainable components with continuing application uptake. Support European capacity to develop and improve AI tools, so that knowledge and feedback generated through HPC software development also strengthen Europe's own software ecosystem.

An Indicative target is for each critical component supported by the programme to have an accountable maintainer team, a public roadmap, an active contributor-development plan and a funded maintenance arrangement extending beyond its initial R&I project.

4.1.5 Data Management I/O and Storage

Establish Interoperability as a First-Class Design Principle

Funding and strategic initiatives should prioritise interoperable interfaces across storage, metadata, and data-management layers, supporting the standardisation of cross-vendor APIs, metadata schemas, semantic contracts, and policy models. Interoperability and portability should be established as key criteria in procurement and project evaluation, alongside incentives for open ecosystems and conformance testing to reduce proprietary vendor lock-in. These measures are essential to enable effective federation, portability, and seamless integration across HPC, AI, and data infrastructures.

Invest in Data-Movement Minimisation Strategies

Prioritise research that avoids data movement rather than optimising each transfer incrementally, supporting compute-to-data, computational storage, in-situ and in-transit processing, and locality-aware scheduling. Develop cross-layer frameworks for data-path and placement optimisation, and fund accountable AI-assisted placement and lifecycle policies that expose uncertainty and permit user correction. Addressing data movement is essential, as it remains a hidden scaling wall connecting performance, energy, congestion, and cost.

Build a Sustainable European Storage and Data Ecosystem

Create targeted, long-term funding mechanisms for European storage technologies, strategic software, and operational capability. Support hardware capabilities including controllers, media, interconnects, packaging, and storage communication. Invest in maintainable software ecosystems spanning file systems, object stores, orchestration, data movers, and data management. Align procurement, open-source sustainability, and innovation funding with long-term ecosystem building. Sovereignty requires sustained capability and viable maintenance paths, not isolated research prototypes.

Make Data Integrity, Provenance, and Observability Measurable

Integrate verification, provenance, and cross-layer observability into HPC and AI data pipelines. Define interoperable standards for integrity checks, lineage, telemetry, and machine-actionable policies. Connect storage and workflow observability to AI explainability and scientific reproducibility. Include integrity, provenance preservation, and explanatory capability in evaluation and benchmarking. Trust in data must be measurable, risk-based, and connected to corrective action.

Develop Standardised Benchmarking for Data Systems

Establish a European benchmarking framework for storage and data-management systems that reflects real workflows. Go beyond peak I/O throughput to include movement efficiency, energy, cost, interference, metadata, integrity, and AI-workload behaviour. Use end-to-end workflows and transparent metric semantics. Support vendor-neutral comparison while allowing system-specific optimisation to be represented fairly. Current benchmarks are too often misaligned with heterogeneous, federated, and AI-era workloads.

Rethink Data Management and I/O Procurement models for Rapid Innovation

Adapt data management and I/O procurement to technologies and workloads that evolve faster than conventional system lifecycles. Enable staged procurement, modular upgrades, and open interfaces. Allow qualified emerging technologies to be integrated after the initial contract. Create protected pathways for experimentation within or alongside production systems. Static procurement is incompatible with rapidly evolving storage, memory, and data-management technologies.

Integrate Sustainability into Storage System-Design Metrics

Make energy and lifecycle cost first-class metrics in storage system design, evaluation, and funding. Fund energy-aware placement, movement, retention, and deletion research. Require lifecycle and cost models in infrastructure and research proposals. Include operational energy, embodied cost, media replacement, and data-retention obligations in benchmarking. Storage is a first-order driver of infrastructure energy use and long-term cost.

4.1.6 Mathematics and Algorithms

The overarching recommendation is to invest in a new generation of mathematical methods, algorithms, libraries and abstractions that deliver portability, reliability, energy efficiency and controllable accuracy across heterogeneous—and increasingly AI- and quantum-augmented and other emerging computing —architectures, thereby also reducing the risk of technology lock-in and strengthening digital sovereignty. This can be divided into the following recommendations:

Scalable Algorithms for Extreme Heterogeneity

Novel mathematical methods, algorithms need to be developed to address heterogeneity and deep hierarchies, and these should enable attaining scalability and efficiency. Thus, mathematical methods should be increasingly tailored to the heterogeneous architecture of supercomputers and made to expose multi-level parallelism and be able to scale to multiple accelerators. The data structures should be adapted and introduce support for heterogeneous accelerators. Communication avoiding and reduced and mixed-precision algorithms should be further developed and become a mainstream to enhance efficiency and scalability. New frameworks that enable automation should be developed to enhance the process.

Communication Avoiding, energy efficient, approximate methods and algorithms

Communication avoiding approaches stemming from domain areas such as Linear Algebra should be generalised since these minimise data movements and enable larger scalability. Mathematical frameworks that enhance approximate computing that spans the High Performance Computing stack, ranging from the hardware level, programming languages and compilers, mathematical algorithms and libraries, up to large scale applications in scientific computing, data analytics and AI Hardware/software must be developed as a co-design approach. All the above approaches combined will enable energy efficiencies.

Hybrid Methods and Algorithms

A fundamental rethink and new mathematical and algorithmic approaches need to be developed to integrate AI stochastic based methods with classical numerical methods. Advanced generalised frameworks are needed to address balancing of errors in the hybrid approach. Future advancements must focus on developing novel strategies to ensure the stability of coupled methods and on establishing robust frameworks for incorporating probabilistic solution distributions into deterministic systems. These efforts are critical to fully leverage AI's potential in scientific computing while preserving the reliability and accuracy of traditional deterministic approaches and bring these to the mainstream.

Resilient and Fault-Tolerant methods and algorithms

Novel mathematical approaches that lead to resilient algorithms that avoid checkpointing and are amenable to dealing with faults such as self-healing algorithms as well as autonomous error detection and correction are really key. Stochastic algorithms allowing oversampling and autocorrection are one of the possible approaches. Enhanced algorithmic fault tolerance approaches should be applied together with novel methods for automatic/silent error correction have to be developed. In algorithm design, failures should be taken into account with the same comprehensive and holistic mindset used by electronic engineers for the design of electronic circuits, and in other engineering disciplines when performing failure mode and effects analysis. There must be quantitative guarantees on performance even in the presence of faults.

Next-Generation Mathematical Methods and Accuracy-Aware, Performance-Portable Libraries

Future research should prioritise the development of adaptive, self-tuning mathematical methods and libraries that deliver high performance, numerical accuracy, reliability and portability across heterogeneous HPC architectures. Particular emphasis should be placed on accuracy-controlled mixed-precision computing, decoupling numerical solution accuracy from the precision used for computation and data storage. Research should advance automated algorithm selection, AI-assisted optimisation,

communication-aware partitioning and dynamic load balancing, while supporting the modernisation of legacy libraries without compromising correctness or reproducibility.

Europe should establish sustainable, open mathematical software ecosystems supported by continuous cross-platform testing, reproducible benchmarking and robust development practices to reduce technology lock-in and ensure long-term portability, maintainability and resilience.

Efficient quantum-classical algorithms

It is recommended to develop an integrated, end-to-end frameworks that treat problem decomposition, classical numerical support, error/resource awareness, execution orchestration, and benchmarking not as separate research threads but as tightly coupled components of a single co-design methodology. This would involve scalable, automated decomposition tools capable of handling arbitrarily large HPC problems by dynamically allocating subproblems to quantum, quantum-inspired, or classical solvers based on real-time resource and noise characteristics, moving beyond today's largely manual or heuristic partitioning schemes.

HPC for AI and AI for HPC

The development of scalable, explainable, and reliable AI methods that are suitable for deployment on future HPC and exascale systems is key. Scientific applications require more than predictive accuracy; they demand interpretability, uncertainty quantification, numerical robustness, reproducibility, and efficient execution across heterogeneous computing architectures.



4.1.7 European Computing Continuum

To unlock the full value of post-exascale infrastructures, Europe should treat the areas and recommendations below as a strategic investment area in their own right. The key objective is not only to deploy more powerful systems, but to ensure that European HPC, cloud, edge, AI, and data infrastructures (The European Computing Continuum) can operate together as an open, trusted, and sustainable continuum for science and innovation.

Build the European continuum as a federated operational ecosystem

Europe should invest in federated mechanisms that allow HPC, cloud, edge, AI, and data infrastructures to interoperate through common interfaces for access, orchestration, metadata, accounting, and operations. The objective is to move from a collection of excellent but fragmented infrastructures to a coherent continuum usable across domains and countries.

Make openness and interoperability mandatory design principles

Future European infrastructures should avoid lock-in in hardware, software stacks, AI frameworks, workflow platforms, and data services. Public investment should systematically favour open standards, portable programming environments, interoperable APIs, and modular software architectures that allow new technologies to be integrated without repeated rewrites of applications and workflows.

Elevate workflow and data orchestration to first-class infrastructure capabilities

End-to-end workflows — not isolated compute jobs — should become the main operational unit of the post-exascale continuum. This requires: interoperable workflow descriptions, common execution APIs, advanced data staging and transfer services, and cross-site orchestration mechanisms able to handle interactive, real-time, and long-running unattended workflows.

Strengthen trusted federation, identity, and sovereignty

A European continuum requires interoperable identity and access management, support for service accounts and automation, secure data sharing across administrative boundaries, and stronger technical foundations for confidentiality, integrity, provenance, and post-quantum readiness. Trust, security, and data sovereignty should be embedded across the whole continuum rather than treated as add-ons.

Invest in software portability, reproducibility, and long-term preservation

Europe should fund portable software stacks, accelerator abstraction layers, reproducible execution environments, FAIR software and workflows, provenance frameworks, and long-term preservation of data, models, workflows, and execution context. These capabilities are essential to maintain scientific value over long lifecycles and to reduce the cost of adapting to new architectures.

Make sustainability a core operational constraint

Future investments should address sustainability at ecosystem level: energy-aware scheduling, minimisation of unnecessary data movement and recomputation, lifecycle assessment of hardware and software, waste-heat reuse, material supply-chain awareness, and reduced duplication across infrastructures. The success of post-exascale systems should be measured not only by performance, but by scientific and industrial value delivered within environmental and societal limits.

Support application readiness and transfer to production.

Many scientific and industrial applications are not yet ready for heterogeneous, AI-enabled, federated execution. Europe should therefore fund application refactoring, co-design with users, reference workflows, portability testbeds, validation platforms, and production-grade software engineering. Bridging the gap between research prototypes and robust operational services should become an explicit R&I objective.

Improve usability and widen access to continuum capabilities

The continuum will not scale if users must manually handle site-specific credentials, queues, data staging, software environments, and hardware choices. Europe should promote high-level interfaces, science gateways, notebook-based services, federated catalogues, and service models such as HPC-as-a-Service that expose capabilities in a secure, portable, and user-oriented manner.

Invest in skills, coordination, and long-term ecosystem stewardship

Sustaining the continuum requires more than technology. Europe should reinforce cross-infrastructure co-design, strengthen the role of research software engineers and data/workflow specialists, and create long-term training and career frameworks aligned with the needs of heterogeneous, AI-enabled, secure, and sustainable infrastructures.

Explore emerging frontiers for the continuum, including HPC/AI in space.

As the computing frontier expands toward airborne, autonomous, and space-based platforms, Europe should anticipate how future ecosystem technologies can federate with these new computational layers. R&I should investigate energy-efficient edge platforms, trusted distributed execution in safety-critical environments, and the integration of data and AI services collected and processed beyond the traditional ground-based continuum.



4.2 Cross cutting technical theme recommendations

4.2.1 Non-Conventional Systems and Architectures

Europe should devote energy and investment to fostering nascent technologies — supporting academic research, investing in start-ups, and funding internal R&D. This strategy rests on the premise that early adoption of innovative technologies confers significant competitive advantages. We provide here more policy-oriented recommendations backing the technical details highlighted in the related separate white paper.

Mitigate the vulnerability of the EU Non-Conventional and Disruptive Technologies ecosystem

The EU ecosystem remains highly vulnerable to export controls, trade restrictions and geopolitical dependencies[61]. This vulnerability applies in particular to nascent and potentially disruptive technology developments. Open-minded foresight and robust strategies are required to reduce these vulnerabilities and strengthen European autonomy in this space. As an example, a civil counterpart of the EU Defence Innovation Scheme (EUDIS[62]) could be established to fund and manage disruptive technologies in the scientific and technical domains selecting promising cutting-edge, high-risk/high-impact technologies leading to game changing impact.

Implement robust, comprehensive programmes to nurture start-ups, incubators, and similar organisations developing Non-Conventional Architectures and Technologies

Such programmes should permit the use of existing EU and EuroHPC JU member states' resources to fund "risky", low-TRL technologies, implement incentives to attract and retain human capital, and actively scout new technologies. The objective is to enable the maturation of these technologies and to prevent their disruption or acquisition by non-EU actors.

Enhance awareness and understanding within EU industry and policy communities of Non-Conventional Architectures and Technologies

The endorsement by policymakers, academia, industries, and enterprises of future Non-Conventional Architectures and Technologies is contingent on a deep understanding of how these technologies work and what their real-world use cases and applications may be. Improving stakeholder awareness and understanding of these fields will enable the selection and support of the most promising development paths and players.

Facilitate the development of European champions and coalitions of stakeholders in academia and industry willing to act as technology advisors and testers of EU Non-Conventional Architectures and Technologies

This should be achieved via targeted programmes, subsidies, and incentives. Such action will also strengthen Europe's position at a time when supply chains are under pressure and competition is intensifying. Forming coalitions of interested parties in Europe will strengthen the strategic posture against heavily funded competitors and protectionist measures by others.

[61] https://digital-strategy.ec.europa.eu/en/library/proposal-chips-act-20
[62] https://eudis.europa.eu/eudis-tracks/disruptive-technologies_en

ETP4HPC could establish a dedicated cross-functional expert group to scout and evaluate Non-Conventional Architectures and Technologies

In its role as a private, industry-led, non-profit association promoting HPC, AI, and quantum research and innovation, and as one of the three private members of the EuroHPC JU represented on the Advisory Boards (RIAG and QTAG), ETP4HPC could assemble a small cross-functional group of subject matter experts (three to five individuals) to liaise with its members and other constituents — including the European Commission, EuroHPC JU and the CHIPS JU — and to interact with start-ups, incubators, and comparable organisations. This group should be tasked with scouting, rapidly evaluating, and validating use cases; identifying showstoppers; and assessing the feasibility, risks, and deployment potential of Non-Conventional Architectures and Technologies.

Establish a dedicated European Modular Experimental Computing Facility for Non-Conventional Architectures and Technologies (EMECF)

The establishment of a European Modular Experimental Computing Facility (EMECF) should be prioritised as a strategic initiative to accelerate the research, validation, and industrial adoption of non-conventional computing architectures for the post-exascale era. Unlike conventional supercomputing installations, the EMECF would serve as a large-scale, modular, and reconfigurable experimental facility for evaluating emerging technologies, including neuromorphic, photonic, in-memory, and quantum-inspired computing, alongside conventional CPU/GPU architectures. Building on Europe's Modular Supercomputing Architecture (MSA), the facility would provide common networking, storage, resource-management, and software infrastructure to enable realistic experimentation with technologies, software stacks, and algorithms. It should incorporate energy-aware design principles and provide open access to European research organisations, industry, start-ups, and SMEs. Such a facility would strengthen European technological sovereignty, reduce dependence on non-European hyperscalers, and accelerate the transition from research prototypes to commercially deployable technologies. ETP4HPC, in collaboration with the European Commission, EuroHPC JU, and Chips JU, could lead the development of a strategic programme and business case for the EMECF within the next 12–18 months.

Establish a structured European maturation and adoption pathway for Non-Conventional Architectures and Technologies

ETP4HPC, working with EuroHPC JU, CHIPS JU, the European Commission, research infrastructures, technology developers, and HPC centres, could connect technology scouting to practical adoption through a common pathway:

Scout → Assess → Match to HPC/AI use cases → Prototype → Benchmark → Integrate → Demonstrate → Productise → Early Adoption.

4.2.2 Quantum Computing

Fault-Tolerant Quantum Computing (FTQC)

Europe should pursue four interconnected priorities in the FTQC domain. Firstly, it should develop a strong QEC ecosystem by advancing scalable decoding, enabling open access interfaces, supporting adaptive and ML-based approaches, and allowing early experimentation on real FTQC hardware; secondly, establishing open integration standards by standardizing interfaces across the quantum–HPC stack, promoting modular architectures, enabling SME participation, and avoiding vendor lock-in through interoperability; thirdly, addressing security and multi-tenancy early by defining secure multi-

user models, implementing isolation mechanisms, controlling access policies, and embedding security across all system layers; finally, building a shared resource estimation framework by creating a common, compiler-integrated layer validated across hardware and QEC approaches to guide FTQC readiness, research priorities, and industrial investments, as well as develop approaches to optimize and measure execution performance as programs scale in size.

Middleware and Scheduling

The following strategic research goals should guide Europe's middleware and scheduling agenda; firstly, enabling controlled multi-tenancy on QPUs by moving beyond exclusive allocation, supporting time-slicing and partitioning, and exposing multi-tenancy concepts through system interfaces but still allowing for user authentication, authorisation, and accounting; secondly, advancing multi-QPU execution as a managed system capability by integrating QPU awareness into HPC resource managers and supporting scalable orchestration across heterogeneous devices; thirdly incorporating device telemetry into scheduling decisions by defining relevant metrics and using calibration-aware placement to improve efficiency and reduce failed executions; finally, establishing stable interfaces and coordinated hybrid scheduling by standardizing interactions between control layers and HPC systems and enabling end-to-end optimization of hybrid quantum–classical workflows.

Workflows, Programming Models, and Benchmarking

Research and development priorities for hybrid workflows should address the following areas; defining standardised workflow design patterns and environments for common hybrid quantum–HPC applications to capture reusable structures and algorithm characteristics across multiple data holders, HPC systems, quantum simulators, and QPU backends; developing abstraction layers for orchestration that decouple application logic from execution details while enabling efficient workflow management; designing robust workflow–middleware interfaces to support job execution, monitoring, error handling, and integration with existing tools and schedulers, and; establishing workflow-level benchmarking frameworks with metrics that capture orchestration efficiency, data movement, and end-to-end performance.

Emulators, Debugging Environments, and Quantum-Inspired Algorithms

The four pillars of emulation, tensor networks, benchmarking, and quantum-inspired methods should be treated as a single co-evolving system with the following priorities; firstly developing high-fidelity emulation ecosystems by calibrating emulators with real hardware, establishing continuous benchmarking feedback loops, and ensuring portability between emulator and device through common interfaces; secondly, extending these emulation ecosystems to expose virtual QPUs integrated within HPC infrastructures, emulate quantum communication networks, and supporting emerging modular and distributed QPU architectures, enabling the development and validation of large-scale quantum workflows; thirdly, scaling tensor network simulations using HPC techniques by investing in distributed (MPI-based) and GPU-accelerated implementations to overcome current memory and performance limitations; finally, accelerating quantum-inspired computing by leveraging quantum concepts to advance classical algorithms and explicitly integrating quantum and classical innovation pathways.

Industrial Applications

Moving from research demonstrations to industrial impact requires coordinated action. A problem-first approach should guide efforts, with early industry involvement to ensure practical relevance and alignment with real-world applications. Shared industrial problem instances and standardized metrics

are needed to track progress and clarify the gap to usable solutions. A clear taxonomy of quantum advantage must be adopted to distinguish different value types and avoid misjudging timelines. Structured international collaboration should be strengthened, building on partnerships such as Q-Neko with sustainability and CO_2 reduction integrated as core objectives. Bridging skills and cultural gaps through multidisciplinary training and co-design is essential to identify and realize genuine quantum advantage. Please also see the next chapter on Industrial Usage (cluster 4 "Quantum readiness for industrial applications").

Training and Education

The following recommendations address the structural gaps in Europe's quantum technologies training landscape (in line with R-4 in the Executive Summary of Recommendations); establishing structured training access to quantum resources by introducing shared teaching credits, education-oriented access calls, dedicated QPU/emulator allocations, and exposure to multiple hardware modalities across the EuroHPC ecosystem; developing role-based, competency-driven training frameworks by tailoring curricula to diverse backgrounds, aligning skills to target roles, and using standardized proficiency models with portable credentialing; and; creating continuous feedback and improvement loops in training by collecting usage data, learner feedback, and adoption metrics, and using them to refine both training programs and underlying quantum technologies



4.2.3 HPC/AI Usage, Applications and Workflows

Build European foundations for portable, scalable applications

The following strategic priorities are recommended to strengthen software portability, adaptability to evolving architectures, and sustained scientific computing performance across European HPC systems.

Make portability a shared European capability by treating interoperable mathematical, communication, and domain-specific libraries as long-lived European infrastructure, with sustained development, maintenance, and performance engineering. Design for changing architectures by using stable, vendor-neutral interfaces with adaptable backends, allowing shared libraries and community codes to exploit new processors and accelerators without repeated application-level rewrites. Preserve leadership-class simulation capability by ensuring systems provide sustained FP64, memory, and interconnect performance alongside AI-oriented capabilities, while supporting precision-aware and hybrid numerical methods where they deliver demonstrated scientific value.

Optimise the ecosystem for scientific productivity

The following priorities are recommended to optimise the European HPC/AI ecosystem for scientific productivity, energy efficiency, and validated scientific outcomes.

Establish reference end-to-end workflows spanning simulation, AI, and hybrid applications, and use measures such as time-to-solution, reproducibility, throughput, and energy-to-solution alongside traditional performance metrics to assess time to validated scientific results. Make energy and scientific productivity part of procurement by evaluating the European HPC/AI ecosystem using representative end-to-end workloads, including the energy required to produce validated results, rather than relying on peak performance alone.

Make data logistics a first-class, interoperable capability

Scientific workflows increasingly depend on data moving between storage, compute, AI services, and external data spaces. Data logistics should therefore be exposed as an integrated capability of HPC and AI environments, rather than left to researchers to manage through site-specific mechanisms. The following priorities are recommended to enable seamless data access, optimise end-to-end workflows, and automate data logistics where appropriate.

Provide transparent, interoperable interfaces across storage tiers, compute resources, and European data spaces, so that applications can access, move, and process data without requiring users to manage the underlying data logistics. Co-design data placement, movement, caching, and processing with workflow execution, supporting in-situ, in-transit, and streaming approaches where they reduce data movement and improve efficiency. Use workflow-aware automation, including AI where appropriate, to optimise placement, prefetching, tiering, and movement while keeping decisions observable, controllable, and reproducible.

Advance AI as a scientific computing capability

The following priorities are recommended to support increasingly capable and efficient scientific AI models while ensuring their scientific validity and reliability.

Invest in scalable training and inference for scientific foundation models and emerging AI approaches. Develop methods and reference workflows that assess accuracy, robustness, uncertainty, and physical consistency, ensuring that AI is evaluated for scientific validity rather than benchmark performance alone.

Build the human capability for the intersection of AI and HPC

The following priorities are recommended to strengthen cross-disciplinary skills in HPC and AI, ensure appropriate human oversight, and promote trustworthy AI practices.

Equip HPC experts to integrate, validate, and supervise AI components and agents, while giving AI researchers the skills to use HPC efficiently, including scaling, scheduling, data movement, and energy efficiency. Define clear boundaries between what AI systems may observe, recommend, and execute, with human oversight for decisions that affect scientific validity, resource allocation, infrastructure operation, or experimental outcomes. Establish common practices for validation, provenance, auditability, and reproducibility, and embed them in training programmes and professional development.



4.2.4 Energy Efficiency and Sustainability

Establish a EuroHPC Experimental Centre for Sustainable Computing.

A dedicated project team, constituted under EuroHPC governance in coordination with ETP4HPC, hosting entities, technology providers, and energy-sector stakeholders, should deliver within 18 months the blueprint of a production-scale facility integrating the full vertical stack of sustainability levers (real-time lifecycle monitoring, energy-aware scheduling, advanced cooling, heat reuse, grid integration, second-life hardware, and sufficiency governance), with a costed implementation roadmap enabling procurement in the first half of the post-2030 decade. Potential synergies with the proposed European Modular Experimental Computing Facility (EMECF) will be explored in a future SRA7 Agile White Paper.

Fund open-source, end-to-end energy measurement tooling (hardware and software) aligned with regulatory reporting.

Computing centres should be equipped with tools that automate the extraction of metrics compliant with the Energy Efficiency Directive (EED), Delegated Regulation 2024/1364, and the Corporate Sustainability Reporting Directive (CSRD), and EuroHPC reporting should be aligned with these instruments to avoid metric fragmentation across the European fleet. This should extend the energy-aware software already under development within the EuroHPC energy-efficiency line rather than open a parallel effort and consolidate its monitoring layer into a minimal telemetry specification — the energy, power, thermal, and utilisation signals every EuroHPC system must expose, with defined granularity, semantics, and access rights — made a mandatory procurement requirement so that systems acquired outside that consortium expose the same interfaces.

Create application-level energy benchmarks and an open European comparison website.

Beyond system-level benchmarks such as HPCG or SPECpower, Europe should establish energy-to-solution and carbon-to-solution benchmarks on representative production workloads, executed under documented and reproducible measurement methodologies, with results published on an open EuroHPC-endorsed web platform, extending the mobilising logic of the Green500 to the application

layer, where most optimisation potential resides. Such a platform is distinct in purpose from the operational data repositories centres maintain for their own tuning: reproducible comparison across sites and architectures, open to non-participating centres and to code developers.

Harmonise sustainability metrics beyond PUE and make them binding.

PUE, WUE, and Energy Reuse Factor - ERF constraints should apply to all new sites; Carbon Usage Effectiveness - CUE, energy-to-solution, carbon-to-solution, and lifecycle impact should be standardized across the EuroHPC fleet; and embodied-carbon disclosure should become a standard procurement evaluation criterion.

Treat energy and carbon as first-class allocation constraints.

Resource and job management systems should incorporate energy and carbon budgets alongside compute, memory, and storage; users should be rewarded for the temporal and spatial flexibility of their jobs; environmental justification should be required for large compute allocations, extending practice already established in some ERC and Horizon Europe instruments, and energy and carbon reporting should become mandatory for publicly funded research and industrial projects.

Activate facility energy assets as grid resources.

Existing multi-megawatt UPS batteries should be operated, within a reserved backup state-of-charge, as active control layers between facility and grid for self-consumption optimisation, peak shaving, and fast frequency response, with dedicated longer-duration storage deployed where multi-hour flexibility or price arbitrage is targeted, under IEEE 1547-2018 and the IEC 62933 series; and waste-heat reuse should be deployed with quantified ERF targets wherever district heating or industrial heat consumers are accessible.

Embed circularity across the EuroHPC fleet.

Procurement frameworks should require design-for-disassembly, repairability, and embodied-carbon disclosure; selective component refresh should replace full-system replacement where practical; and cascaded second-life hardware (training → inference → academic computing → education) should be validated at HPC scale following the Reduce–Reuse–Repair–Remanufacture–Recycle–Recovery hierarchy.

Make eco-design literacy a core competency of the European HPC community.

A mandatory eco-design and sustainable computing module should be integrated into the EUMaster4HPC curriculum (in line with R-4 in the Executive Summary of recommendations), complemented by continuing-education tracks for system administrators delivered through EuroHPC competence centres, and by developer-facing feedback loops linking code changes to measured energy outcomes.

5 Industrial Usage: Priorities for European HPC, AI and Data Ecosystems



We summarise in this chapter the R&I trends and challenges for European HPC as AI based tools, methods and techniques have proliferated in the last few years with the advent and usage of GPGPUs and the continued development of the AI software stacks with exploitation of increased volumes of data. More details corresponding to each of these R&I trends and challenges can be obtained in the individual White Papers for the Research Domains and the Thematic Trends. These need to be recognised and addressed especially from a policy perspective.

5.1 Industrial context and strategic objective

High-performance computing (HPC), artificial intelligence (AI), data-intensive engineering and emerging computing technologies are becoming inseparable from industrial competitiveness. Simulation remains fundamental to the design and validation of complex products, while AI is increasingly used to optimise processes, interpret operational data and support engineering decisions. At the same time, industrial computing is becoming more heterogeneous and distributed, spanning large computing centres, cloud platforms, company data centres and edge systems close to production assets.

In Europe, these developments offer major opportunities to reduce development cycles, widen design exploration, improve quality, lower resource consumption and create new services. Yet access to computing capacity alone does not translate into industrial value. Adoption is still constrained by fragmented software and data environments, difficult integration with established engineering processes, security and sovereignty concerns, uncertain costs, a shortage of specialised skills and weak pathways from pilot projects to operational deployment. These barriers are particularly significant for SMEs, but they also affect large enterprises that depend on long-lived validated applications, commercial software and complex supply chains.

Exemplified by sectors such as European Aerospace and Defence, and European Automative, our central recommendation is therefore to move from isolated technology capabilities towards integrated, trustworthy and accessible industrial computing environments. Research and innovation must be application-driven and assessed not only by technical performance, but by whether solutions can be deployed, governed and maintained within real industrial processes. This requires coordinated progress across eleven closely connected areas.

Some overlaps exist between these recommendations and those in the previous chapter, reflecting the perspectives and priorities of the independent working groups, recognising the cross-cutting nature of the challenges, and highlighting areas where coordinated action across multiple research and technology domains is necessary.

Strategic research and innovation priorities are provided below, with 5.2.1, 5.2.5 and 5.2.11 having a higher priority to be addressed compared to the others.

5.2 Strategic research and innovation priorities

5.2.1 Sovereign industrial AI and HPC convergence

HPC and AI are increasingly combined in simulation, optimisation, generative design, analytics and operational decision-making. However, industrial users still encounter separate infrastructures, incompatible data formats and fragmented specialist teams. Europe should develop integrated environments that support physics-informed AI, surrogate models, AI-assisted simulation and industrial foundation models while protecting proprietary data, models and design knowledge. Priority should

be given to hybrid HPC–AI programming models, efficient learning from disparate data sources, scalable inference, confidential execution and lifecycle governance. The European Commission and EuroHPC JU, working with industry, research organisations and European Digital Innovation Hubs, should prioritise these capabilities within their R&I programmes from 2027 onwards, leveraging AI Factories and European HPC infrastructures to support industrial demonstrators. These demonstrators should demonstrate measurable gains in accuracy, engineering productivity, energy efficiency and time to solution, with credible pathways from experimentation to production. This priority is in line with R1 – Strengthen European AI Sovereignty and Accelerate AI-Driven Scientific Discovery, and R3 – Give European SMEs and Industry Easier Access to Advanced Computing (from the Executive Summary of recommendations).

Recommendation

Develop integrated, secure HPC-AI environments around validated industrial use cases and reusable reference workflows. Prioritise interoperability, model governance and protection of industrial data so that solutions can move from pilots to routine engineering operations.

5.2.2 Industrial digital twins and real-time simulation

Industrial digital twins are evolving from static models into persistent systems linking simulation, operational data and AI throughout an asset's lifecycle. Their potential includes predictive maintenance, anomaly detection, operational optimisation and faster feedback into product design. Realising this potential requires scalable multi-physics models, real-time data assimilation, uncertainty management and trustworthy hybrid AI–physics methods. Research should address the full operational loop across edge, cloud and HPC resources. Long-running demonstrators should evaluate accuracy, latency, maintainability, resilience and business value under realistic operating conditions, rather than stopping at offline proofs of concept.

Recommendation

Establish interoperable digital-twin architectures that connect validated models, operational data and lifecycle systems across the computing continuum. Fund long-running industrial demonstrators that prove accuracy, maintainability and operational value beyond isolated proofs of concept.

5.2.3 Emerging computing architectures for industrial workloads

Heterogeneous systems combining CPUs, GPUs, AI accelerators, advanced memory and domain-specific processors can improve performance and energy efficiency, but they also disrupt industrial software stacks developed and validated over decades. The priority is not merely access to new hardware, but a sustainable migration path for existing applications. Research should advance performance-portable programming models, compilers, runtime systems, memory-aware algorithms, automated code transformation and mature debugging and profiling tools. Hardware–software co-design and benchmarking must use representative industrial workflows and include development effort, migration cost, energy efficiency and long-term maintainability—not peak performance alone.

Recommendation

Couple investment in emerging hardware with sustained programmes for application modernisation, developer tools and performance-portable software. Evaluate architectures against representative industrial workflows, including migration cost, energy efficiency and long-term maintainability.

5.2.4 Quantum readiness for industrial applications

Quantum computing may eventually benefit optimisation, materials modelling, chemistry, risk analysis and selected simulation problems, but industrial value remains uncertain and application-dependent. Near-term activity should focus on readiness through hybrid quantum–classical workflows, realistic emulation, portable software environments and transparent comparison with the strongest classical methods. Research should support hybrid orchestration, error mitigation, resource estimation, shared benchmarking and application discovery. Industrial pilots should publish negative as well as positive findings and define credible technical and economic readiness criteria, allowing companies to build capability without premature commitment or platform lock-in.

Recommendation

Create application-led quantum-readiness programmes that compare hybrid quantum approaches with the strongest classical alternatives. Support portable toolchains, emulation and shared benchmarks so companies can build expertise without committing prematurely to a single technology.

5.2.5 Trusted, secure and sovereign industrial computing

Trust and control over sensitive assets are prerequisites for industrial use of shared infrastructures. Industrial workloads can include proprietary geometries, manufacturing data, software source code, AI models and information subject to regulation or export restrictions. Security must therefore extend from hardware and networks to identities, containers, workflows, registries and data spaces. Priorities include confidential computing at scale, zero-trust architectures, remote attestation, secure multiparty workflows, verifiable software and AI supply chains, fine-grained usage controls and resilient recovery. Common assurance frameworks and independent validation are needed so that industrial users can evaluate risks consistently across providers and federated environments.

Recommendation

Embed security, confidentiality and policy enforcement into computing services by design rather than adding them after deployment. Establish common assurance frameworks that allow industrial users to evaluate trust, compliance and operational risk across federated environments.

5.2.6 Sustainable and energy-aware HPC and AI

The energy demands of simulation and AI are growing while industry faces pressure to reduce costs, emissions and exposure to limited power availability. Hardware efficiency alone is insufficient: total impact also depends on workload design, software quality, utilisation, cooling, data movement and electricity mix. Energy and carbon should become visible optimisation variables alongside time, accuracy and cost. Research should advance comparable application-level metrics, energy- and carbon-aware scheduling, efficient training and inference, mixed precision, data reduction, holistic monitoring and lifecycle assessment. Demonstrators should quantify impacts on design quality, throughput and total cost of ownership.

Recommendation

Adopt standard, application-level metrics that relate energy and carbon impact to useful industrial outcomes. Make energy-aware scheduling, software optimisation and lifecycle assessment standard capabilities of future HPC and AI services.

5.2.7 Industrial workflow automation and autonomous engineering

Industrial workflows combine geometry preparation, meshing, simulation, optimisation, data processing, AI development and reporting, yet often rely on manual handovers, bespoke scripts and scarce expert knowledge. Controlled automation can shorten cycle times, improve reproducibility and expand design exploration. Research should focus on composable workflow descriptions, policy-aware orchestration, automated verification, reusable engineering agents, and integration between MLOps and simulation operations. In safety-critical and regulated applications, human-in-the-loop designs must specify where automation may act and where expert approval is mandatory. Open interfaces, provenance and validation of AI-generated actions are essential to prevent opacity and platform lock-in.

Recommendation

Develop open, traceable workflow frameworks that integrate simulation, AI and data operations while retaining human control over critical decisions. Validate automation through measurable reductions in cycle time, failure rates and dependence on bespoke scripts.

5.2.8 A federated HPC–cloud–edge continuum

Industrial data and workloads increasingly move between operational edge devices, enterprise platforms, cloud services and large HPC systems. Today these environments often use different identity, scheduling, monitoring and data-transfer mechanisms, forcing workflows to be redesigned for each deployment. Europe should develop a coherent operating model based on federated identity, portable execution, predictable data movement, common observability and policy-driven resource placement. Orchestration should consider latency, confidentiality, cost, energy and performance. End-to-end demonstrators must test realistic recovery, security and service-level responsibilities, not only technical connectivity.

Recommendation

Build federated services around common identity, orchestration, observability and data-movement mechanisms. Demonstrate complete industrial workflows across edge, cloud and HPC resources with clear governance, service levels and failure responsibilities.

5.2.9 Portable software and resilient software ecosystems

Industrial HPC depends on commercial packages, internal codes, open-source libraries and long-lived validated workflows. Modern architectures cannot achieve industrial impact if required applications are unavailable, unaffordable or impractical to migrate. Software strategy must address performance and deployment portability, reproducible builds, automated testing, sustainable maintenance and licensing. Coordinated evaluation licences, usage-based commercial models and European-level agreements could allow companies—especially SMEs—to test applications on new platforms while providing vendors with evidence of demand. Open-source alternatives also require long-term maintenance, documentation and industrial validation.

Recommendation

Combine software-modernisation programmes with practical licensing mechanisms that let industrial users evaluate commercial applications on advanced systems. Require portability, reproducibility and sustainable maintenance plans for software receiving shared investment or strategic support.

5.2.10 Industrial data spaces and semantic interoperability

Industrial AI, digital twins and collaborative engineering depend on data that can be discovered, interpreted and reused across organisational and technical boundaries. The challenge is not simply moving data, but preserving its meaning, provenance, quality and permitted use. Research should develop governed data spaces, adaptable sector-level semantics, digital threads, scalable knowledge graphs, federated discovery and machine-readable usage controls. Solutions must work with legacy systems and real supply chains; a single universal ontology is unlikely to be practical. Demonstrators should measure reductions in data preparation, reconciliation errors and duplicated work while proving that intellectual property and contractual rights remain protected.

Recommendation

Develop sector-aligned semantic frameworks and governed data spaces that preserve provenance, meaning and usage rights across the industrial lifecycle. Test them in real value chains and measure reductions in data preparation, integration errors and duplicated effort.

5.2.11 Industrial skills, adoption and competence ecosystems

Infrastructure access and training courses alone do not create durable industrial capability. Companies need people who combine domain knowledge with HPC, AI, data engineering, cybersecurity and software deployment. This need is particularly acute for SMEs and organisations trying to progress beyond an initial proof of concept. Programmes should provide embedded expert teams, shared expert pools, research collaboration and repeatable onboarding from use-case definition through deployment and internal ownership. Support should also cover organisational change, procurement, cybersecurity and business-case development, and be accessible to companies with limited proposal-writing capacity.

Recommendation

Provide accessible, embedded expert support that helps companies progress from use-case selection to operational deployment and internal ownership. Evaluate adoption programmes by sustained use, skills retained and measurable business outcomes rather than participation counts alone.

5.3 Measuring industrial impact

Progress should be evaluated through application-level outcomes rather than infrastructure statistics alone. Relevant indicators include the proportion of industrial workflows progressing from pilot to routine operation; reductions in time to solution, engineering cycle time and data-preparation effort; workflow failure rates and reproducibility; energy and carbon intensity per useful engineering outcome; the availability and portability of required commercial and open-source software; and the proportion of supported companies retaining skills and continuing operational use after external assistance ends.

European industry does not primarily lack computing capacity. It lacks the integrated, secure, interoperable and usable environments required to turn that capacity into engineering and business outcomes. Infrastructure investment must therefore be matched by sustained investment in software, security, data, energy efficiency, workflow integration and people. Doing so will allow European HPC, AI and emerging computing technologies to deliver measurable industrial value while reinforcing competitiveness, sovereignty and resilience.

6 European HPC Technology SMEs



6.1 State of play

Europe enters the post-exascale period with a strong shared HPC infrastructure, expanding AI Factories and Antennas, semiconductor initiatives and a growing strategic-autonomy agenda, but the position of its HPC Technology SMEs continues to remain uneven[63]. These SMEs contribute across processors, accelerators, storage, cooling, interconnects, middleware, system integration, AI/HPC workflows, various services and expertise/consulting, and domain applications, and their specialisation allows them to translate research into deployable products and fill strategically important niches. Nevertheless, major procurements, consortium structures and financing models continue to favour large integrators and established multinational vendors, leaving SMEs technically visible but often weak in budget share, decision authority, customer ownership, intellectual-property control and routes to market. The central state-of-play issue is therefore not whether SMEs participate, but whether Europe gives them the capital, test environments, procurement access and strategic influence needed to own, scale and export sovereign HPC technologies.

We next list our strategic recommendations for European HPC Technology SMEs. We also list recommendations that also applies to us/ETP4HPC as an industry association.

6.2 Strategic recommendations

6.2.1 Recognise HPC technology SMEs as a strategic capability class

Recommendation: EuroHPC JU, the European Commission and Member States should explicitly distinguish HPC Technology SMEs from general SME digital adopters. Policy instruments should recognise their role as technology owners, supply-chain anchors and post-exascale capability builders.

This distinction makes supply-side technology creation a specific policy objective. It allows programmes to assess SMEs by the European capabilities, intellectual property and supply-chain resilience they create, rather than treating them only as users of digital infrastructure.

6.2.2 Create SME-first tracks in European programmes and initiatives

Recommendation: Dedicated calls should support SME-led R&I in processors, accelerators, RISC-V and Arm ecosystems, chiplets, EDA workflows, cooling, interconnect, storage, AI/HPC software stacks, system integration, workflow tools and domain applications. Calls should be sized and administered for SME delivery capacity.

SME-first tracks would align call size, administration and leadership expectations with SME capacity. They would also direct investment toward strategically important parts of the HPC and AI stack where specialist European companies can develop defensible products.

6.2.3 Reform procurement to move SMEs from subcontractors to prime contractors

Recommendation: Public procurements should use “lots” or subdivisions of larger contracts, pre-commercial procurement, innovation partnerships and framework models that allow SMEs to lead

[63] https://etp4hpc.eu/wp-content/uploads/wpfd/preview_files/ETP4HPC_WP_SME-ecosystem_20260115%28f0467e856f9ee5f05a1815fca47b7787%29.pdf

where they hold specialist technology. Tenders should include minimum SME participation by budget value and decision authority, not merely number of participants!

Procurement is the bridge from publicly funded innovation to durable market position. Smaller lots and innovation-oriented instruments give SMEs a realistic route to customer ownership, while value-based participation requirements reduce token involvement in large consortia which is happening today.

6.2.4 Bridge the prototype-to-product gap for European HPC technology SMEs

Recommendation: Europe should establish coordinated funding and industrialisation pathways to help HPC Technology SMEs bridge the gap between validated research prototypes and commercially deployable products, particularly across Technology Readiness Levels (TRLs) 5–8. Support should cover engineering maturation, access to EDA tools and test chips, certification, security hardening, documentation, support models, reference deployments, market entry and first-of-a-kind industrial pilots.

Promising research does not become a dependable commercial product without substantial investment in engineering, qualification, customer validation and business development. This transition is particularly challenging for HPC Technology SMEs, which may lack the financial resources and industrial partnerships needed to progress beyond publicly funded R&I projects. The European Commission, EuroHPC JU and Chips JU should coordinate existing funding instruments, including Chips JU pilot lines and design infrastructure for semiconductor technologies, EIC Transition and Accelerator programmes for technology maturation and commercialisation, EuroHPC-supported validation and demonstration activities, and complementary national innovation and industrialisation funds. These mechanisms should provide a clearer pathway from research outcomes to commercially viable European HPC products, strengthening European intellectual-property ownership, industrial competitiveness and technology sovereignty.

6.2.5 Build sovereign AI/HPC software and middleware around SMEs

Recommendation: SME-led software should be prioritised in compilers, runtimes, scheduling, workflow orchestration, containerisation, data movement, observability, energy optimisation, AI model pipelines, and inference services in European federated HPC/AI infrastructures.

Prioritising SME-led software strengthens European control over the layers that determine portability, resource use and workflow execution. It also creates opportunities for European companies to own reusable platforms rather than contribute isolated components. It is to be noted that Europe already has a very rich ecosystem of SME developed HPC software and middleware.

6.2.6 Make AI Factories and Gigafactories SME product engines

Recommendation: AI Factories should provide staged SME support from discovery through product validation and procurement readiness. Gigafactory governance and procurement should include SME technology integration, open interfaces and measurable SME value share from inception.

The objective is to move beyond just access to compute for these SMEs. A staged pathway from experimentation through eventual procurement would make these infrastructures active launch environments for European SME products and services.

6.2.7 Establish a European HPC SME catalogue and coordination hub

Recommendation: A curated catalogue should map out SME capabilities, TRLs, reference deployments, technical domains, procurement readiness and contact points. A coordination hub should connect SMEs to EuroHPC access, AI Factories, NCCs, testbeds, procurement opportunities and investors.

A common capability map would address fragmentation and make credible SME suppliers easier to discover. The coordination function would then turn visibility into practical links with infrastructure, finance, partners, customers and procurement opportunities.

6.2.8 Create SME-integrated talent pathways

Recommendation: EUMaster4HPC, NCCs, Centres of Excellence, doctoral networks and AI Factory training programmes should include SME placements, industrial challenge projects, secondments, technical sales training and shared expert pools for SMEs.

Embedding SMEs in education and mobility programmes would improve their access to scarce technical and commercial skills. It would also expose students and researchers to product engineering and industrial deployment, not only academic HPC practice.

6.2.9 Establish an SME policy and political advocacy group

Recommendation: ETP4HPC could help create a formal SME Policy and Political Advocacy Group - a lobbying and policy-engagement function - to promote the interests of HPC Technology SMEs with the European Parliament, European Commission, EuroHPC JU, RIAG, INFRAG, QTAG and Member States. Its role should be to address the perception and risk that procurement and strategy are disproportionately shaped by large integrators and non-European incumbents, and to ensure that SME evidence, product capability and market barriers are visible in strategic and programme level decision making.

A permanent advocacy function would give HPC Technology SMEs a coordinated voice in policy, programme and procurement decisions. Evidence gathered from SMEs could then be presented consistently before rules and roadmaps become fixed.

6.2.10 Measure what matters

Recommendation: EuroHPC, ETP4HPC and relevant programmes could publish annual metrics on SME participation by budget value, contract value, prime-contractor roles, IP ownership, reference deployments, product outcomes, follow-on revenue, cross-border scaling and contribution to European supply-chain resilience.

Outcome based indicators would show whether SME policy is producing ownership, revenue, products and strategic influence. Annual reporting would also enable programme and procurement rules to be corrected when participation remains superficial.

6.3 Conclusion

European HPC Technology SMEs are essential for converting Europe's research strength, public infrastructure and post-exascale ambition into sovereign products, resilient supply chains and industrial value. Infrastructure access alone will not achieve this: SMEs need programme tracks designed for their

capacity, funding across the full prototype-to-product pathway, routes into real procurement value, open and interoperable markets, access to talent and a sustained policy voice. The above recommendations translate this objective into practical support through capability mapping, coordination, one-stop-shop access, testbeds, procurement tools, targeted finance, skills pathways and accountable governance. Europe's success should ultimately be judged by whether its SMEs gain intellectual-property ownership, reference deployments, customers, revenue and strategic influence, rather than simply appearing as minor participants in projects and tenders.



7 Conclusion and Acknowledgements

This ETP4HPC Strategic Research Agenda 7 (SRA7) Master, released in late 2026, sets out Europe's key Research and Innovation priorities for HPC, AI and quantum computing over the next five years. It is supported by 13 thematic White Papers that provide the technical background, challenges and detailed recommendations underpinning this strategic vision.

Looking ahead, four priorities stand out: strengthening European technological sovereignty and industrial competitiveness; building an interoperable, federated and sustainable computing ecosystem; advancing European hardware, software and emerging computing architectures; and accelerating the convergence of HPC, AI and quantum computing. Progress in these areas will require coordinated investments, stronger links between research and industrial adoption, and sustained collaboration across the European computing ecosystem.

The next steps are to gather feedback from the European HPC community, finalise and release the supporting White Papers, and align SRA7 priorities with the EuroHPC Multi-Annual Strategic Programme (MASP) and Europe's emerging post-exascale roadmap. ETP4HPC encourages policymakers, funding organisations, research institutions, technology providers and industrial users to engage with these recommendations and translate them into coordinated R&I programmes, investments and concrete actions that strengthen Europe's long-term competitiveness and technological autonomy.

The supporting white papers are as follows:

Research Domains

- System Architecture
- System Hardware Components
- System Software and Management
- Data Management, IO & Storage
- European Computing Continuum
- Programming Environment
- Mathematics and Algorithms

Cross cutting technical thematic trends

- Energy Efficiency and Sustainability
- Quantum Computing
- Non-Conventional Architectures
- HPC/AI Applications and Workflow

Policy Topics

- SMEs
- Industrial Usage

Additional agile White papers supporting the SRA7 will be released in the next two years based on sub-topics that warrant deeper discussion.

7.1 Acknowledgements

The ETP4HPC would like to thank the SRA7 working group leaders below for providing valuable inputs that contributed to this Master. The working group leaders have worked with a team of around 130 European experts providing invaluable contributions over the last year.

- Alberto Scionti (LINKS Foundation)
- Andreas Wierse (SICOS)
- Andy Forrester (HypeAccelerator)
- Craig Prunty (SiPearl)
- Diego Andrade Canosa (UDC)
- Fabrizio Magugliani (E4 Computer Engineering)
- Gabriel Antoniu (Inria)
- Johannes Gebert (HLRS)
- Luca Carriere (Avio Aero)
- Manolis Marazakis (FORTH)
- Manuel Torres Rodrigues (SURF)
- Marc Duranton (CEA)
- Matteo Barborini (University of Luxembourg)
- Michael Sleigh (ECMWF)
- Nico Mittenzwey (MEGWARE)
- Paul Carpenter (BSC)
- Pierre Seroul (Bull)
- Robert Haas (IBM)
- Sagar Dolas (SURF)
- Sarah Neuwirth (JGU Mainz)
- Utz-Uwe Haus (HPE)
- Valeria Bartsch (Fraunhofer)
- Valery Guilleaume (Nodeum)
- Vassil Alexandrov (STFC)

This work is edited by the ETP4HPC Office with the following team:

- Dr. Sai Narasimhamurthy (Chief Editor)
- Pascale Bernier Bruna
- Marcin Ostacz
- Gabriella Povero
- Dominique Grelet



ETP4HPC.EU

Cite as: Narasimhamurthy S. et al. ETP4HPC's SRA 7 – Strategic Agenda for High-Performance Computing in Europe – 2026. ETP4HPC https://doi.org/10.5281/zenodo.24202154

DOI: 10.5281/zenodo.24202154

© ETP4HPC 2026