ATOMIX – Where Technology and Vision Converge
Gammel Havnevej 5
4070 Kirke Hyllinge
T: +45 26176088
E: contact@atomix.dk
W: atomix.dk

European Scientific High-Performance Computing and the Mandate for Digital Sovereignty in the Age of AI

A Technical Guide for Enterprise IT Leaders

Executive Summary

The European Union (EU) is executing a coordinated strategic response to address critical global technology dependencies by merging its massive, continent-spanning High-Performance Computing (HPC) investment with the urgent infrastructure demands of generative Artificial Intelligence (AI). The EuroHPC Joint Undertaking (JU) acts as the central instrument for this strategic pivot. By leveraging its established exascale foundation—including flagship systems like JUPITER, LUMI, and Leonardo—the EuroHPC JU is rapidly deploying specialized AI Factories (AIFs) across the continent. These facilities are fundamentally designed not merely as centers of raw computational throughput, but as highly compliant, digitally sovereign ecosystems.

The AI Factory model ensures control over the entire technological value chain, spanning from indigenous hardware development through the European Processor Initiative (EPI) and the creation of secure, regulated data spaces compliant with the GDPR and the AI Act, to the implementation of advanced operational pipelines via MLOps and extensive skills development through National Competence Centres (NCCs). The unprecedented deployment velocity, evidenced by the selection of 19 AI Factories and supporting Antennas between 2024 and 2025, underscores Europe's commitment to achieving robust technological autonomy. This systemic investment positions Europe as a global leader in the development of *trustworthy*, *human-centric* AI, ensuring that computational capacity is strategically secured and aligned with core European values and regulatory principles.

ATOMIX – Where Technology and Vision Converge
Gammel Havnevej 5
4070 Kirke Hyllinge
T: +45 26176088
E: contact@atomix.dk
W: atomix.dk

1. The Strategic Nexus of HPC, AI, and Digital Sovereignty

1.1 Defining Digital Sovereignty in the Age of Exascale

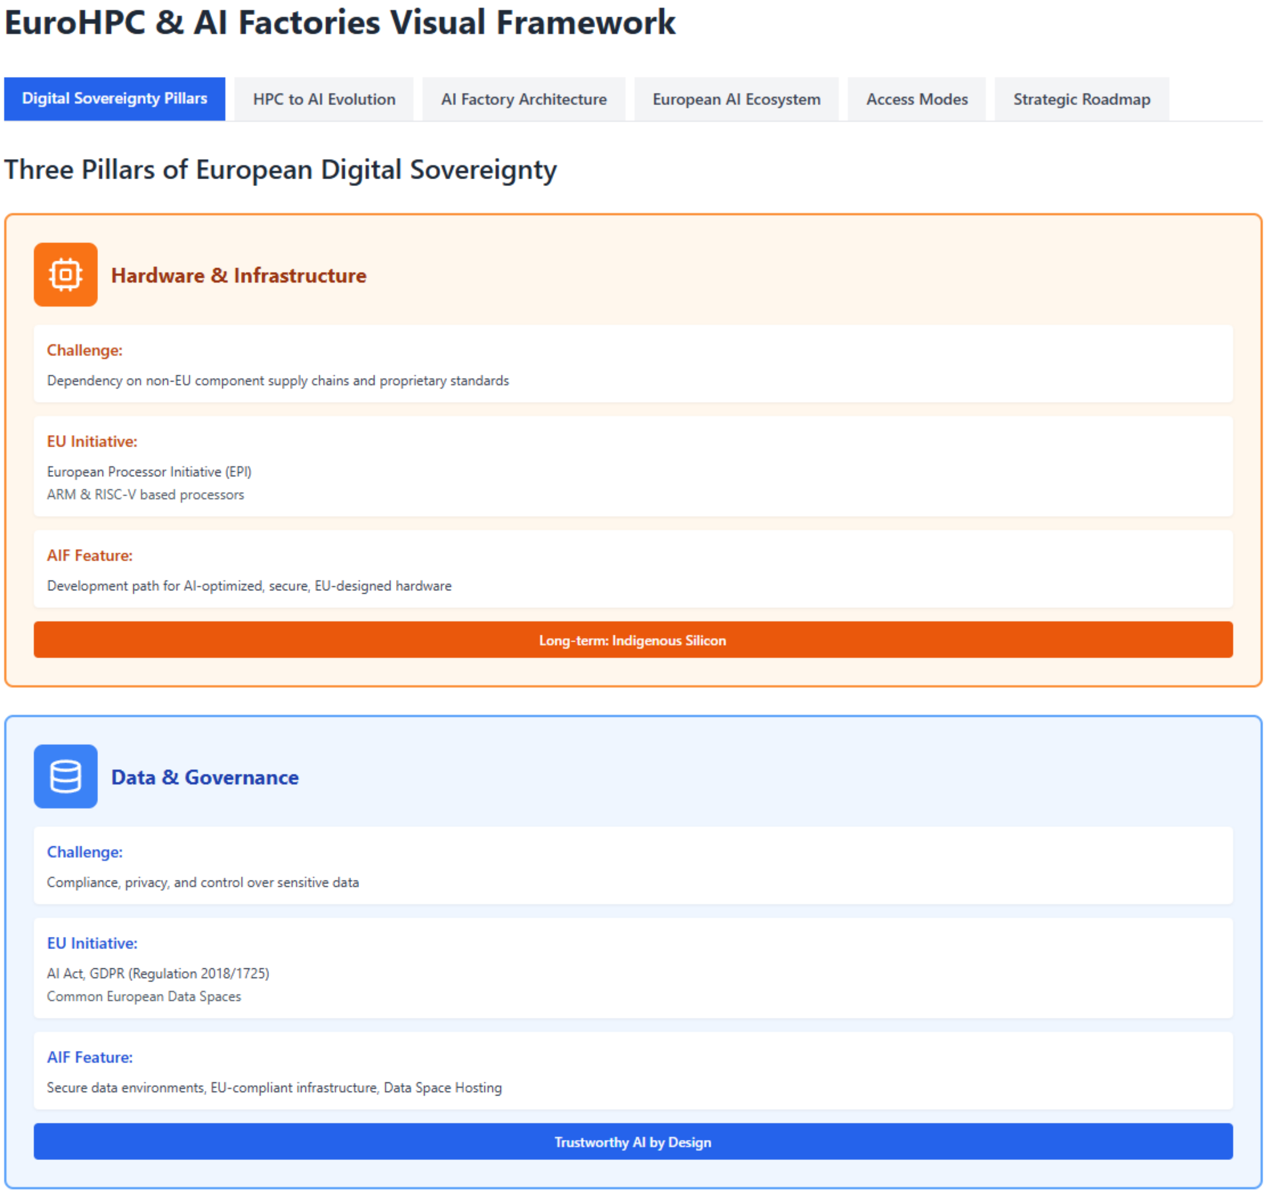

Digital sovereignty is defined as the foundational capability of a state or union to control its own digital destiny, encompassing governance over essential pillars such as data, hardware, and software.[1] For the European Union, the EuroHPC JU's core mission is explicitly linked to supporting the digital transformation of its economy while ensuring the technological sovereignty of the continent.[2] This goes beyond mere access to computing power; it involves deep control over the underlying infrastructure and technological supply chains.

High-Performance Computing is recognized as a critical geopolitical infrastructure, underpinning foundational national capabilities. Supercomputing systems are vital for enabling breakthrough scientific discoveries, driving complex industrial innovation across sectors such as aerospace and pharmaceuticals, and reinforcing national security, defense, and sovereignty. This includes utilizing massive computational resources in the fight against cyber-criminality and for the robust protection of critical national infrastructures.[2] The strategic importance of HPC thus transcends commercial utility, becoming an integral component of national and continental resilience. Furthermore, the JU strategy is structurally aligned with the EU's overarching political goals, supporting the twin green and digital transitions and contributing directly to the specific digital goals enshrined in the "Digital Decade" vision for 2030.[3]

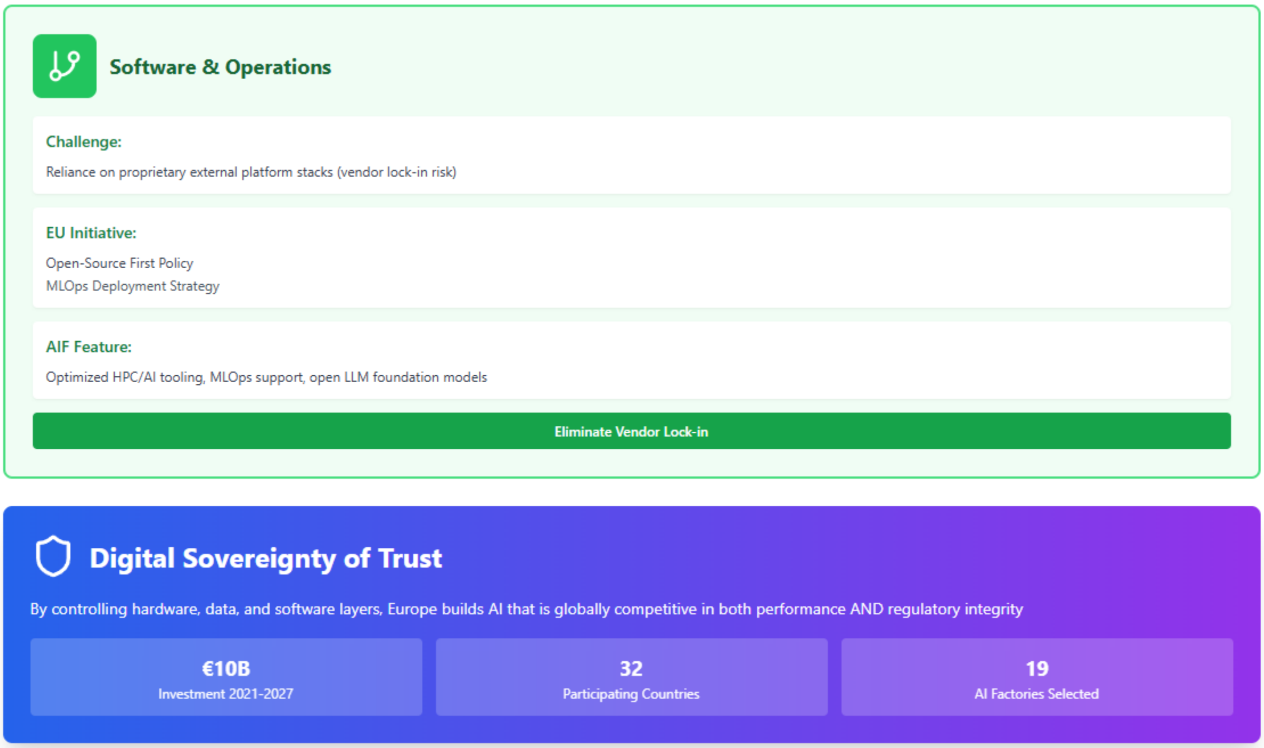

1.2 EuroHPC JU: A Decisive Geopolitical and Economic Instrument

The EuroHPC JU represents a landmark initiative designed to overcome the scale limitations faced by individual Member States. It is established as a legal and funding entity that successfully pools the resources of the European Union, 32 participating European countries, and several private partners.[4] The scale of this coordinated investment is immense: overall EU and Member State investments dedicated to supercomputing infrastructures and the emerging AI Factories are projected to reach €10 billion during the 2021–2027 period.[5]

The core objective of the JU is to develop, deploy, extend, and maintain a world-leading federated, secure, and hyper-connected ecosystem encompassing supercomputing, quantum computing, service, and data infrastructure across the European territory.[2] This vision necessitates the development of novel applications optimized for these systems and based on a supply chain that guarantees the availability of components, technologies, and knowledge, thereby actively limiting the risk of global disruptions.[2] Critically, no single European country possesses the capacity or financial means to develop world-class supercomputing resources individually. The Joint Undertaking provides the essential financial support, coordination mechanism, and risk-sharing framework necessary to achieve globally competitive scale.[4] The federation model ensures that these resources, once deployed, are accessible to a wide range of public and private users throughout Europe.[4]

1.3 The Foundational Shift: From Traditional Simulation to AI Workloads

The integration of AI into the EuroHPC framework necessitates a fundamental acknowledgment of the architectural disparity between traditional HPC and modern Machine Learning (ML) workloads. Traditional HPC, often optimized for large-scale physical simulations, is typically CPU-bound. In stark contrast, training deep learning models—the foundational activity for modern AI—is computationally intense, often consuming 100-1000 times more computational resources than conventional applications during training phases.[6] These ML workloads depend critically on specialized Graphics Processing Unit (GPU) acceleration, shifting the architectural bottleneck away from CPUs and onto parallel processing capabilities.[6]

This architectural shift imposes unique and unprecedented demands across the entire infrastructure stack. For compute, there is a hard reliance on specialized GPU architectures, such as the NVIDIA Tesla V100, A100, or H100 series, which represent current enterprise standards optimized for parallel mathematical operations.[6] For data, AI systems process datasets measured in petabytes. This necessitates storage architectures optimized for high-throughput, read-heavy access patterns, unlike transactional database architectures. Model training requires storage systems capable of sustained throughput rates reaching 10 to 50 GB/s, whereas standard enterprise storage arrays often deliver only 1-5 GB/s, creating severe performance bottlenecks.[6] Furthermore, distributed AI training requires sophisticated, high-bandwidth, low-latency networking, often relying on specialized interconnects such as NVIDIA NVLink or InfiniBand to ensure efficient gradient synchronization across nodes.[6]

The cost of infrastructural inadequacy in this domain translates directly into a serious competitive and geopolitical risk. Organizations attempting to deploy advanced AI on unprepared infrastructure commonly experience extended training times—extending model training from days to months—which severely impacts development velocity and time-to-market for AI-driven features.[6] Furthermore, insufficient memory or storage throughput forces developers to resort to smaller datasets or reduced model complexity, directly compromising the resulting model's accuracy and commercial value. The investment in AI Factories, therefore, is an explicit and urgent strategic measure taken by Europe to mitigate this competitive risk and prevent external infrastructural limitations from compromising continental AI development capabilities.[6]

ATOMIX – Where Technology and Vision Converge
Gammel Havnevej 5
4070 Kirke Hyllinge
T: +45 26176088
E: contact@atomix.dk
W: atomix.dk

EuroHPC & AI Factories Visual Framework

Digital Sovereignty Pillars | **HPC to AI Evolution** | AI Factory Architecture | European AI Ecosystem | Access Modes | Strategic Roadmap

Evolution: Traditional HPC → AI-Optimized Infrastructure

Traditional HPC (Pre-AI Era)

Compute Focus
CPU-bound operations for physical simulations

Workload Type
Large-scale scientific simulations, weather modeling, molecular dynamics

Architecture
Optimized for numerical precision and deterministic computation

Primary Users
Academic researchers, national labs, scientific institutions

AI-Optimized HPC (AI Factory Era)

Compute Focus
GPU-accelerated parallel processing
100-1000x more intensive during training

Workload Type
Training LLMs with trillions of parameters, generative AI models

Architecture
Specialized GPUs (A100/H100), high-bandwidth memory, NVMe storage (20+ GB/s)

Primary Users
AI startups, SMEs, researchers + traditional scientific users

Continued on the next page…

⚠ Critical Infrastructure Shifts Required for AI

Compute
Traditional: CPU-focused
AI Era: 100,000+ GPUs needed

Storage
Traditional: 1-5 GB/s
AI Era: 20-50 GB/s NVMe

Memory
Traditional: 64-128 GB/node
AI Era: 256-512 GB/node

Network
Traditional: Standard Ethernet
AI Era: InfiniBand/100GbE

EuroHPC Evolution Timeline

2021-2023	**Exascale Foundation** JUPITER, LUMI, Leonardo deployed - Traditional HPC focus
2024	**AI Pivot & Regulatory Amendment** Council amends EuroHPC regulation - AI Factories mandate established
2024-2025	**Rapid AI Factory Deployment** 19 AI Factories selected across 15 Member States + Antennas
2025+	**Integrated AI Ecosystem Operational** 100,000+ GPUs, sovereign LLM training, quantum integration

ATOMIX – Where Technology and Vision Converge
Gammel Havnevej 5
4070 Kirke Hyllinge
T: +45 26176088
E: contact@atomix.dk
W: atomix.dk

2. The Exascale Foundation: Securing Computational Capacity

2.1 Exascale Infrastructure Deployment: JUPITER, LUMI, and Leonardo

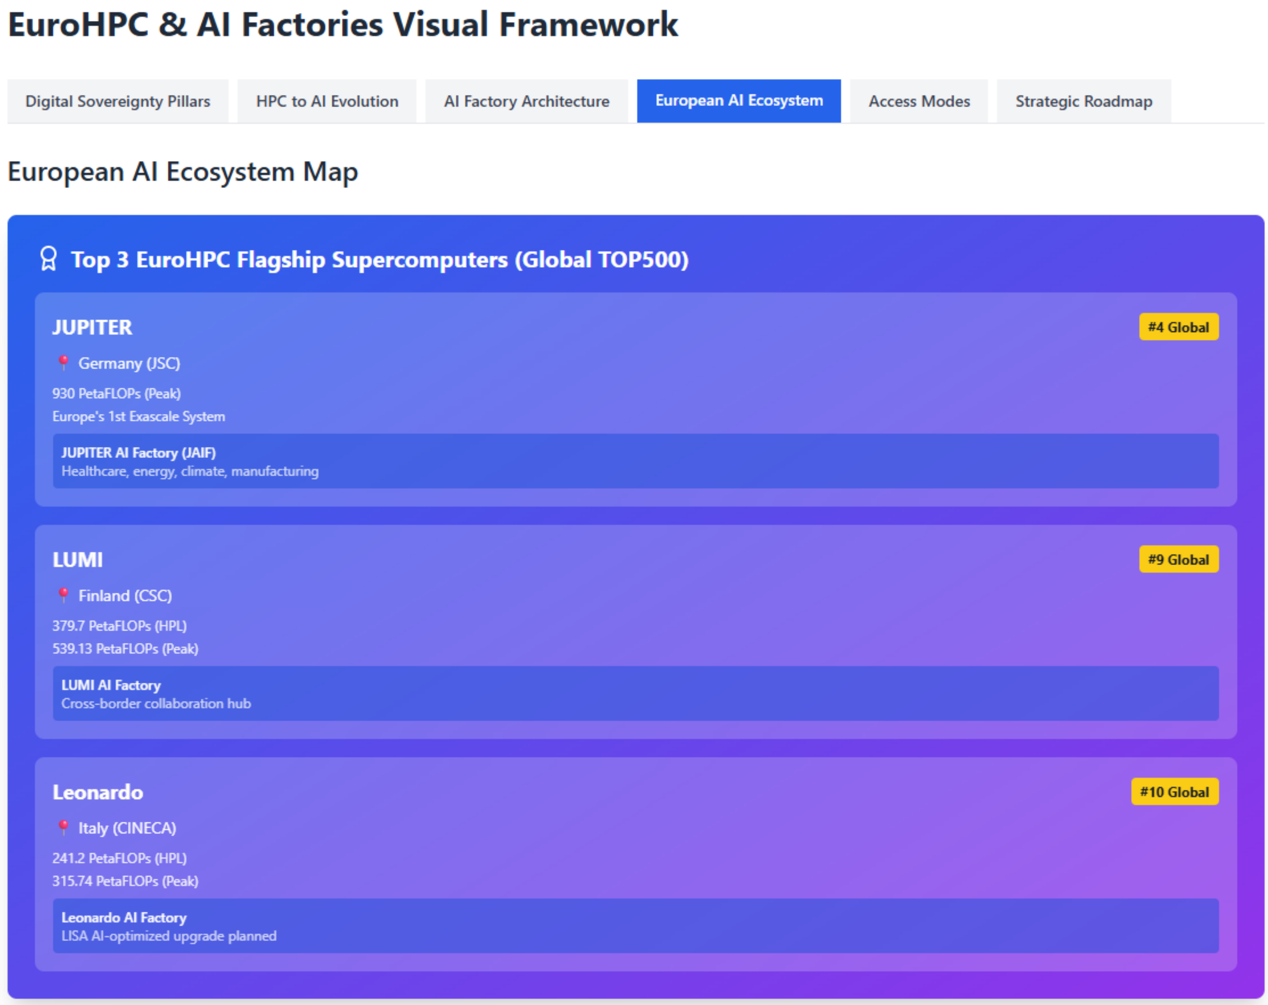

Europe’s strategy to secure global leadership is founded on acquiring and operating world-class computational resources. This is evidenced by the successful procurement and deployment of several flagship systems under the EuroHPC JU. As of June 2025, three of these EuroHPC supercomputers are ranked among the world's top 10 most powerful systems: JUPITER, LUMI, and Leonardo.[7] Securing these top-tier global rankings is crucial for European scientists and industrial users, who can access these systems via the EuroHPC access modes.[7]

JUPITER, Europe's first Exascale supercomputer, hosted in Germany, represents a definitive milestone in this effort. It will soon be capable of delivering 1 ExaFLOP of computing power.[9] Currently ranked fourth globally, JUPITER is based on BullSequana XH3000 direct liquid cooled architecture supplied by Eviden.[10] Beyond its sheer performance, JUPITER has been designed with an emphasis on sustainability, utilizing the waste heat generated during operation to heat buildings and integrating with the Jülich campus heating network.[10] The other major systems, LUMI in Finland (ranked 9th with 379.7 petaflops HPL performance) and Leonardo in Italy (ranked 10th with 241.2 petaflops HPL performance), further cement Europe's position at the forefront of global supercomputing.[9]

These foundational systems are being immediately integrated into the new AI mandate through direct optimization. JUPITER is complemented by the JUPITER AI Factory (JAIF), selected in March 2025, specifically designed to help industry, startups, and SMEs harness the exascale power for the development of secure, data-protection-compliant AI applications.[10] Similarly, Leonardo is undergoing a major upgrade with a new AI-optimised partition (the LISA upgrade), which will enhance its operational performance and provide the necessary functionalities to address evolving user needs in the AI domain.[10]

Table 1: EuroHPC Flagship Supercomputers: Performance and AI Factory Status (As of Q4 2025)

System	**Location**	**Measured Performance (Petaflops)**	**Global TOP500 Rank (June 2025)**	**AI Factory Integration**
JUPITER	Germany (JSC)	930.00 (Peak PFLOPs)	#4	JUPITER AI Factory (JAIF) [10, 11]
LUMI	Finland (CSC)	379.7 (HPL), 539.13 (Peak)	#9	LUMI AI Factory [8]
Leonardo	Italy (CINECA)	241.2 (HPL), 315.74 (Peak)	#10	Leonardo AI Factory (AI-optimised upgrade planned) [10]

2.2 The European Processor Initiative (EPI) and Hardware Sovereignty

A critical element of the EU's Digital Sovereignty strategy is addressing the hardware layer, where the majority of global dependency currently resides. The European Processor Initiative (EPI) is established as one of the cornerstones of this strategic plan.[12] The project is instrumental in developing European supercomputing technologies capable of competing effectively on the global HPC market and ensuring that the key competence of high-end chip design remains within Europe.[12]

The EPI focuses on the development of competitive low-power microprocessor units and accelerators, specifically targeting the incoming European Exascale machines. This effort is concentrated on processor technologies based on the ARM and RISC-V Instruction Set Architectures (ISA).[13] The long-term success of the EPI is essential for the economic sustainability and security of European supercomputing infrastructure. The initiative involves finalizing the development of the first generation of low-power microprocessor units and accelerators, enhancing existing technologies, and developing the second generation while ensuring clear paths for industrialization and commercialization, extending to adjacent markets such as edge computing (e.g., autonomous shuttles or video surveillance).[13]

However, the current architecture of the AI Factories presents a calculated strategic compromise regarding hardware autonomy. While the long-term goal of digital sovereignty necessitates indigenous European hardware (EPI), the immediate and urgent need to build AI capacity relies heavily on integrating world-leading commercial GPU accelerators, such as the NVIDIA A100 or Hopper GPUs, which are currently necessary for reaching competitive performance levels.[8] This approach acknowledges that achieving technological sovereignty is a staged process: Europe must accept short-term reliance on existing global supply chains to rapidly deploy the capacity required to compete in the current AI race, while simultaneously funding the EPI to systematically eliminate that reliance in future generations of European HPC and edge computing systems. The failure of the EPI to deliver competitive hardware components would severely compromise the hardware component of Europe’s long-term digital sovereignty strategy.

---

3. The Accelerated Emergence of European AI Factories (AIFs)

3.1 Policy Catalysis: The Regulatory Amendment and Political Mandate

The establishment of the AI Factories represents a politically decisive pivot within the EuroHPC mandate. The transition was rapidly catalyzed by the European Commission, with political guidelines issued in 2024 prioritizing immediate access to new, tailored supercomputing capacity for AI start-ups and industry.[14] This urgency was formalized when the Council of the EU amended the Regulation of the EuroHPC Joint Undertaking in 2024, granting the public-private initiative the legal framework necessary to establish and manage the AI Factories across Europe.[14]

This regulatory amendment was part of the EU's broader AI innovation package. Its main objective was to simplify and accelerate access for public and private users, particularly European start-ups and SMEs, to the expensive, large-scale resources—specifically AI-optimised supercomputers—required for compute-intensive tasks like the training and development of general-purpose AI models.[14] The deployment velocity following this amendment has been rapid and expansive. Following initial calls for Expressions of Interest, the EuroHPC JU selected 7 initial consortia in December 2024, followed by 6 more in March 2025, and another 6 in October 2025, resulting in a total of 19 selected AI Factory sites across Europe, spanning 15 Member States and 2 EuroHPC Participating States.[5]

3.2 The AIF Mandate: Bridging HPC Capacity with Generative AI Demand

The core mandate of the AI Factories is focused squarely on developing trustworthy, cutting-edge generative AI models. These facilities are designed as large-scale computing centers dedicated to the development and training of next-generation AI models containing trillions of parameters.[5] To achieve this demanding goal, AI Gigafactories must bring together massive computing power—projected to be over 100,000 advanced AI processors—alongside a strong emphasis on guaranteed power capacity, reliable supply chains, advanced networking, and operational energy efficiency.[5]

The essential difference between a general HPC center and an AI Factory lies not only in raw computational power but also in its *integration* and *compliance*. They are fundamentally designed to offer customized AI compute capacity on infrastructure that is secure, scalable, and explicitly compliant with EU regulations.[15] This structural compliance ensures that AI models are developed within a trusted, legally defined European ecosystem. The Factories are explicitly contrasted with generic private cloud services, acting as subsidized, strategic capacity tailored for innovation and sovereignty.[8]

AIFs are establishing specialized sectoral foci across Europe to address key industrial and societal challenges. For example, the Dutch AI Factory (NLAIF) puts special emphasis on AI applications involving highly sensitive data, such as health information or IP-protected datasets, focusing on sectors like healthcare, safety, and agriculture.[15] Similarly, the JUPITER AI Factory (JAIF) concentrates on strategic sectors including healthcare, energy, climate change, and manufacturing.[11]

3.3 Cross-Border Collaboration and Ecosystem Building

The operational design of the AI Factory network emphasizes collaboration and accessibility across the continent. Each AI Factory is mandated to function as a "one-stop shop" nationally, providing a single point of contact for European AI startups, SMEs, and researchers.[11] These shops provide comprehensive support, including access to AI-optimised HPC resources, training, technical expertise, and services to develop AI-ready data.[15]

The network is inherently collaborative and designed for cross-border cooperation. For instance, the Gaia AI Factory in Poland will actively collaborate with the PIAST and LUMI AI Factories to create an integrated AI ecosystem across Central Europe.[15] This collaboration is further expanded through the selection of 'AI Factory Antennas.' These Antennas, such as those established in Iceland (AIFA ICE) and Ireland (AIF IRL-Antenna), link their national AI ecosystems and users to the massive compute resources of the larger, central AI Factories (such as LUMI, JUPITER, and the French AIF).[16] This federation model, where smaller local entities connect to the massive centralized compute pool, ensures that the benefits of strategic European investment reach a wider range of public and private users across all participating nations.[4]

ATOMIX – Where Technology and Vision Converge
Gammel Havnevej 5
4070 Kirke Hyllinge
T: +45 26176088
E: contact@atomix.dk
W: atomix.dk

19 AI Factories Across Europe (2024-2025 Selections)

Primary AI Factories (15 Member States)

1. **JUPITER AI Factory (JAIF)**
   Germany - Healthcare, Energy, Climate, Manufacturing
2. **LUMI AI Factory**
   Finland - Cross-border collaboration
3. **Dutch AI Factory (NLAIF)**
   Netherlands - Healthcare, Safety, Agriculture (sensitive data)
4. **Gaia AI Factory**
   Poland - Central European ecosystem

+ 15 additional AI Factories across EU Member States

AI Factory Antennas (Satellite Access Points)

AIFA ICE (Iceland)
Connected to LUMI AI Factory

AIF IRL-Antenna (Ireland)
Connected to JUPITER & French AIF

Federation Model Benefits:
- Local ecosystem access point
- Links to massive central compute
- Distributed benefits across all nations
- Single point of contact nationally

National Competence Centres (NCCs)

Essential gateways bridging the skill gap for HPC & AI adoption

Core Functions:
- Scientific & technical expertise
- Consulting services for SMEs
- Comprehensive training programs
- Map of HPC competences across EU

Supporting: EVITA Training Academy

Data Infrastructure Integration

Seamless connection with broader European data strategies

Connected Platforms:
- Common European Data Spaces
- EOSC (European Open Science Cloud)
- GAIA-X (industrial/commercial)
- Secure data space hosting services

€200M Connecting Europe Facility

Ecosystem Integration: "One-Stop Shop" Model

Each AI Factory serves as a single national point of contact providing:

- AI-optimized HPC Access
- Training & Expertise
- AI-Ready Data Services
- Compliance Support

4. Architectural Blueprint and Technical Specifications for AI Sovereignty

4.1 Compute Architecture: Specialized Infrastructure for LLMs

The construction of the AI Factories demands an infrastructure fundamentally different from previous HPC generations, requiring specialized architectures optimized for deep learning and Large Language Models (LLMs). The technical assessment framework highlights that training these massive models requires substantial dedicated GPU infrastructure. Flagship EuroHPC systems, such as those utilized by AIFs, are equipped with hundreds or thousands of high-performance accelerators, featuring GPUs like the NVIDIA A100 Tensor Core with significant High Bandwidth Memory (HBM2) capacity (40 GB or 80 GB per GPU).[8] Large language models require substantial RAM allocation, with 256-512 GB per training node becoming the standard for serious ML operations.[6]

Beyond the GPU itself, the surrounding CPU architecture is critical for managing data preprocessing, orchestration, and inference serving. Optimal configurations typically allocate 8 to 16 CPU cores per GPU to ensure that CPU resources do not create a bottleneck for data feeding the powerful GPU accelerators.[6] The highly experimental nature of projects like the Horizon Europe Project TrustLLM, which aims to align LLMs with the European vision of human-centric AI, underscores the necessity of this advanced infrastructure. Training different variants of LLMs relevant to this project requires substantial GPU resources to process trillions of tokens, and the infrastructure must be robust enough to support scaling experiments aimed at both increasing model capabilities and decreasing model sizes for better energy efficiency.[17]

ATOMIX – Where Technology and Vision Converge
Gammel Havnevej 5
4070 Kirke Hyllinge
T: +45 26176088
E: contact@atomix.dk
W: atomix.dk

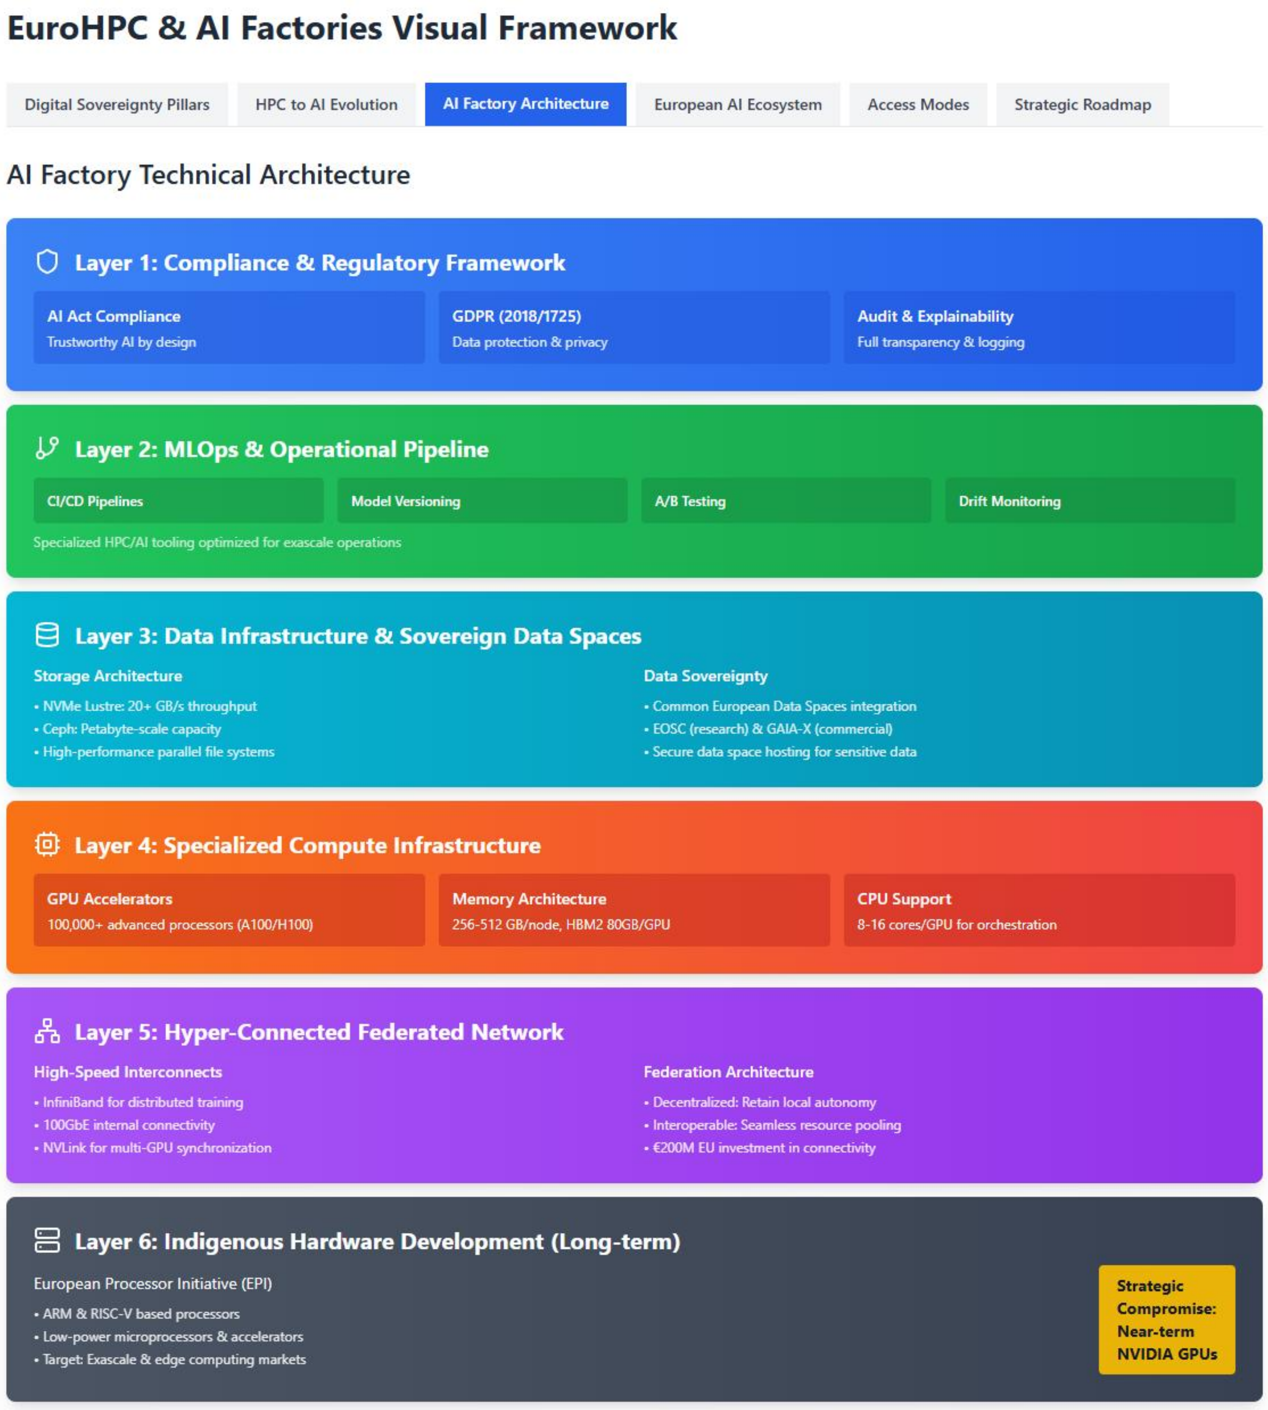

4.2 Data Infrastructure and Management for Sovereign Data Spaces

AI workloads place extreme demands on data infrastructure, necessitating a complete overhaul of conventional storage approaches. The continuous reading of large datasets during training operations requires High-Performance Storage (HPS) systems. This performance requirement is met by implementing NVMe SSD arrays or parallel file systems capable of achieving 20+ GB/s of sustained throughput.[6] Examples from existing EuroHPC systems show deployment of high-performance NVMe Lustre partitions (e.g., 1PB) alongside massive, scalable storage like Ceph (e.g., 23PB), designed to handle the velocity and volume requirements of AI and High-Performance Data Analytics (HPDA).[9]

Crucially, the data infrastructures associated with EuroHPC are mandated to connect seamlessly with broader European data strategies. This includes integration with the Common European Data Spaces, the European Open Science Cloud (EOSC) for academic research, and GAIA-X, which focuses on industrial and commercial services.[4] This strategic federation allows for coordinated and managed data reuse across the EU, supporting secure service provisioning to a wide range of public and private users.[19] To further solidify data sovereignty, the AI Factories are specifically designed to include "data space hosting" services, ensuring that organizations dealing with sensitive data (such as in healthcare or finance) can develop, test, and scale their AI solutions entirely within a secure, EU-regulated, and sovereign environment.[15]

4.3 Networking and Federation: Ensuring Hyper-Connectivity and Resilience

Distributed AI training relies entirely on high-bandwidth, low-latency networking between compute nodes, often requiring 25GbE/100GbE capabilities or specialized AI networking solutions like InfiniBand.[6] Recognizing the paramount importance of connectivity, the EU allocated €200 million through the Connecting Europe Facility to improve the interconnection of HPC, quantum, and data resources, ensuring linkage with the Union's common European data spaces and secure cloud infrastructures.[4]

The concept of a Federated Computing Infrastructure (FCI) is foundational to EuroHPC's architectural design, directly contributing to resilience, scalability, and sovereignty.[21] This federation mandates decentralization, meaning each participating organization retains autonomy over its resources, ensuring that no single external regulatory or geographical jurisdiction can compromise the entire network.[21] This minimizes the impact of external disruptive events, as resources are distributed across locations and organizations. Furthermore, interoperability is designed into the FCI to ensure that different systems, built on varying technologies and standards, can function seamlessly, pooling computational power for the wider European user base.[21] This sophisticated hyper-connectivity is the architectural underpinning of continental computational autonomy.

4.4 Implementing MLOps for Sovereign AI Lifecycle Management

The EuroHPC JU recognizes that simply providing raw compute power is insufficient; the operational efficiency of AI workloads must be addressed systematically. The shift to cutting-edge Machine Learning and LLMs introduces significant new requirements for efficient operationalization on HPC systems.[18]

ATOMIX – Where Technology and Vision Converge
Gammel Havnevej 5
4070 Kirke Hyllinge
T: +45 26176088
E: contact@atomix.dk
W: atomix.dk

A key strategic action outlined by the JU is the explicit mandate to "Deploy and advance HPC workflow tools for AI and HPDA, including MLOps".[18] This strategy addresses the full range of operational challenges identified in the initial whitepaper template, specifically covering the need for continuous integration/continuous deployment (CI/CD) pipelines, sophisticated model lifecycle management (including versioning and validation), and advanced monitoring for performance and drift.[6] AI Factories are required to provide access to EuroHPC systems with specific HPC/AI tooling and advanced expertise in using efficient HPC/AI methods at scale.[18] By integrating these operational best practices—which are currently dominated by global tech platforms—into a controlled, compliant European framework, the JU ensures that the development of sovereign AI is sustainable, scalable, and operationally mature.

---

5. Governing Trusted AI: Sovereignty in Data, Regulation, and Software

5.1 Data Sovereignty Requirements and Secure Processing

For AI Factories to support strategic sectors like healthcare and finance, uncompromising data protection is essential. The Factories commit to offering customized capacity on infrastructure that is explicitly compliant with EU regulations and provides secure, dedicated environments for highly sensitive training data and intellectual property.[15]

As a body of the European Union, the EuroHPC JU is legally obliged to comply with Regulation (EU) 2018/1725, which governs the processing of personal data by Union institutions.[22] This ensures adherence to strict data protection principles, including data minimization (collecting only necessary data), purpose limitation (using data only for specified purposes), and data accuracy.[22] Beyond personal data, the Joint Undertaking is required to ensure the protection of all sensitive information whose disclosure could damage the interests of its members or participants, ensuring commercial and strategic confidentiality within the AI development process.[19] This rigorous application of European law to the computational infrastructure is central to establishing trust in the resultant AI models.

5.2 The AI Act and the Mandate for Trustworthy AI

The AI Act (Regulation (EU) 2024/1689) is the world’s first comprehensive legal framework on AI, designed to foster trustworthy AI and position Europe as a global leader in responsible technology.[23] The AI Factories are integral components of the EU's wider AI Innovation Package, working in conjunction with the AI Act to guarantee safety, fundamental rights, and human-centric AI.[23]

The deep connection between the AI Act and the AI Factories represents a powerful strategic convergence: by controlling the foundational infrastructure, the EU proactively embeds regulatory adherence into the very environment where models are trained. This means European-developed models, trained on compliant AIF infrastructure using secure EU data spaces, inherently meet the requirements for auditability, explainability, and risk mitigation required by the AI Act.[6] This strategic control creates a competitive advantage for EU-developed models globally, based on a "Digital Sovereignty of Trust," where the technology's integrity is guaranteed by the development environment's regulatory alignment. To support rapid adoption, initiatives like the forthcoming AI Act Service Desk are being established to provide tailored compliance guidance, ensuring that regulatory requirements are streamlined and accessible for SMEs leveraging the AIF resources.[24]

5.3 Software Autonomy: Fostering Open and Optimized Ecosystems

Digital sovereignty extends fundamentally to the software layer, requiring strategic measures to reduce dependency on proprietary, external platform stacks which introduce the risk of vendor lock-in. Europe's policy action emphasizes strengthening the open software ecosystem through initiatives that promote, and potentially mandate, an 'open-source first' approach.[25]

The EuroHPC JU actively supports this strategic direction through its application pillar, promoting the development of an innovative and competitive European HPC ecosystem through co-design with users and exploiting synergies with AI and Big Data technologies.[3] Projects utilizing EuroHPC infrastructure directly contribute to software autonomy. For instance, the TrustLLM project focuses on scaling experiments for foundation models that are explicitly "openly shared with the EU LLM community".[17] This commitment to open sharing of foundational knowledge prevents the emergence of proprietary monopolies over essential AI building blocks, ensuring that European industry and academia retain control over the core algorithms and models powering the digital economy. Furthermore, the JU provides support for Centres of Excellence in HPC applications and industrialization of HPC software, ensuring that novel algorithms, codes, and tools are optimized for future generations of supercomputers.[4]

Table 2: Strategic Alignment: EuroHPC Initiatives Supporting Digital Sovereignty

Sovereignty Pillar	**Strategic Challenge**	**EuroHPC/EU Initiative**	**Relevant AIF Feature**
Hardware & Infrastructure	Dependency on non-EU component supply chains and proprietary standards	European Processor Initiative (EPI) [12]	Development path for AI-Optimized, secure, EU-designed hardware [13]

Data & Governance	Compliance, privacy, and control over sensitive data (GDPR/AI Act)	AI Act [23], GDPR (Regulation 2018/1725) [22]	Secure data environments, EU-regulations compliant infrastructure, Data Space Hosting [15]
Software & Operational Stack	Reliance on proprietary, external platform stacks (lock-in risk)	Open-Source First Policy [25], MLOps Deployment Strategy [18]	Provision of optimized HPC/AI tooling, MLOps support, and open LLM foundation models [17]

6. Operationalizing Access and Fostering the Ecosystem

6.1 Access Modes for Industrial Innovation (SMEs and Startups)

EuroHPC & AI Factories Visual Framework

Digital Sovereignty Pillars | HPC to AI Evolution | AI Factory Architecture | European AI Ecosystem | **Access Modes** | Strategic Roadmap

AI Factory Access Modes for Innovation

Free Access for EU SMEs & Startups

Democratizing access to expensive GPU-accelerated infrastructure to foster widespread AI adoption across the EU economy

0€	3	19
Cost for Innovation Access	Access Tiers Available	Factory Locations

Playground Access 1

Compute Resources
Limited resources for initial experimentation

Target Users
Entry-level AI users, startups new to HPC, SMEs exploring AI

Purpose
Initial experimentation and exploration with low barrier to entry

Entry Point
Learn & Experiment

Fast Lane Access 2

Compute Resources
Up to 50,000 GPU hours
Scalable for medium-scale projects

Target Users
AI SMEs already familiar with HPC environments

Purpose
Medium-scale innovation, prototyping, accelerated development

Development Scale
Prototype & Validate

The EuroHPC JU has implemented a sophisticated, tiered access system designed to democratize access to these critical resources, actively addressing the financial barrier that often prevents small entities from engaging with expensive, GPU-accelerated infrastructure. This democratization is vital for fostering widespread AI adoption across the EU economy and ensuring that the benefits of strategic computational power are dispersed.[26] For AI SMEs and startups, access for innovation purposes is open and free-of-charge.[27]

Large Scale Access

3

Compute Resources
Over 50,000 GPU hours
Massive capacity for demanding workloads

Target Users
Advanced users with demanding AI workloads, LLM developers

Purpose
Training large language models (LLMs), industrial scaling, foundation models

Production Scale
Train & Deploy LLMs

Access Mode Progression

Playground	Limited Resources	Entry-level users
Fast Lane	≤ 50,000 GPU Hours	HPC-familiar SMEs
Large Scale	> 50,000 GPU Hours	Advanced LLM training

Important Operational Note
While AI Factories provide immense strategic value with free, subsidized access, they operate on allocated capacity rather than on-demand elastic scaling. For unlimited on-demand capacity, large private cloud vendors may be considered for premium workloads, though without sovereignty guarantees.

Strategic Access Benefits

For SMEs & Startups:
✓ Zero-cost access to world-class infrastructure
✓ Level playing field with large corporations
✓ Accelerated AI innovation cycles
✓ Expert support and training included

For Europe:
✓ Widespread AI capability distribution
✓ Reduced dependency on external platforms
✓ Competitive EU AI ecosystem
✓ Sovereign, trustworthy AI development

The Industrial Innovation Access track is structured to accommodate varying levels of technical expertise and computational need:

- **Playground Access:** This mode provides limited resources, serving as an entry point for users new to AI or HPC, enabling initial experimentation with a low barrier to entry.[27]

- **Fast Lane Access:** Targeted at AI SMEs already familiar with HPC environments, this mode offers scalable access up to 50,000 GPU hours for medium-scale innovation and accelerated prototyping.[27]
- **Large Scale Access:** This tier is reserved for the most demanding AI models and applications, catering to requirements exceeding 50,000 GPU hours, necessary for training and industrial scaling of large foundation models.[27]

While this subsidized model provides immense strategic value and rapid availability, it is important to note the operational distinction: AI Factories, while invaluable, cannot entirely match the on-demand, elastic scaling capacity provided by large private cloud compute vendors, where capacity limits can often be bypassed for a premium.[8]

Table 3: AI Factory Industrial Innovation Access Modes and Target Users

Access Mode	Compute Resources Allocation	Target User/Maturity Level	Purpose
Playground Access	Limited Resources	Entry-level AI users (SMEs, startups)	Initial experimentation and exploration; low barrier to entry [27]
Fast Lane Access	Up to 50,000 GPU Hours	AI SMEs familiar with HPC	Medium-scale innovation, prototyping, and accelerated development [27]
Large Scale Access	Over 50,000 GPU Hours	Advanced users, demanding AI workloads	Training large language models (LLMs) and industrial scaling [27]

6.2 The Role of National Competence Centres (NCCs) in Skill Development

Digital sovereignty is equally dependent on human capital and skill development. The National Competence Centres (NCCs) are critical mechanisms designed to bridge the skill gap and facilitate the uptake of HPC and AI applications by SMEs, academia, and public administrations.[26]

The NCCs function as essential gateways, consolidating and providing scientific and technical expertise, consulting services, and comprehensive training opportunities across Europe.[26] Their mission includes creating a transparent map of HPC competences and institutions and promoting the use of cutting-edge computational technologies to enable businesses to innovate and compete.[29] Furthermore, the JU supports centralized training initiatives, such as the EuroHPC-funded EVITA project (EuroHPC Virtual Training Academy), which aims to build a unified, high-quality training

ATOMIX – Where Technology and Vision Converge
Gammel Havnevej 5
4070 Kirke Hyllinge
T: +45 26176088
E: contact@atomix.dk
W: atomix.dk

framework for emerging technologies, ensuring that the necessary HPC and AI skills are developed and readily accessible throughout European industry and science.[2]

6.3 Future Trajectories: Quantum Integration and Hybrid Cloud Strategy

The long-term vision for the European compute ecosystem extends beyond Exascale and current AI workloads. The EuroHPC JU is mandated to strengthen Europe's leadership in quantum computing technologies, accelerating the development of industrial-scale quantum platforms and integrating these resources into the federated HPC ecosystem.[19] The JU has launched calls, such as the Quantum Grand Challenge Call, to support these initiatives, ensuring a forward-looking strategy that addresses the next frontier of computation.[7]

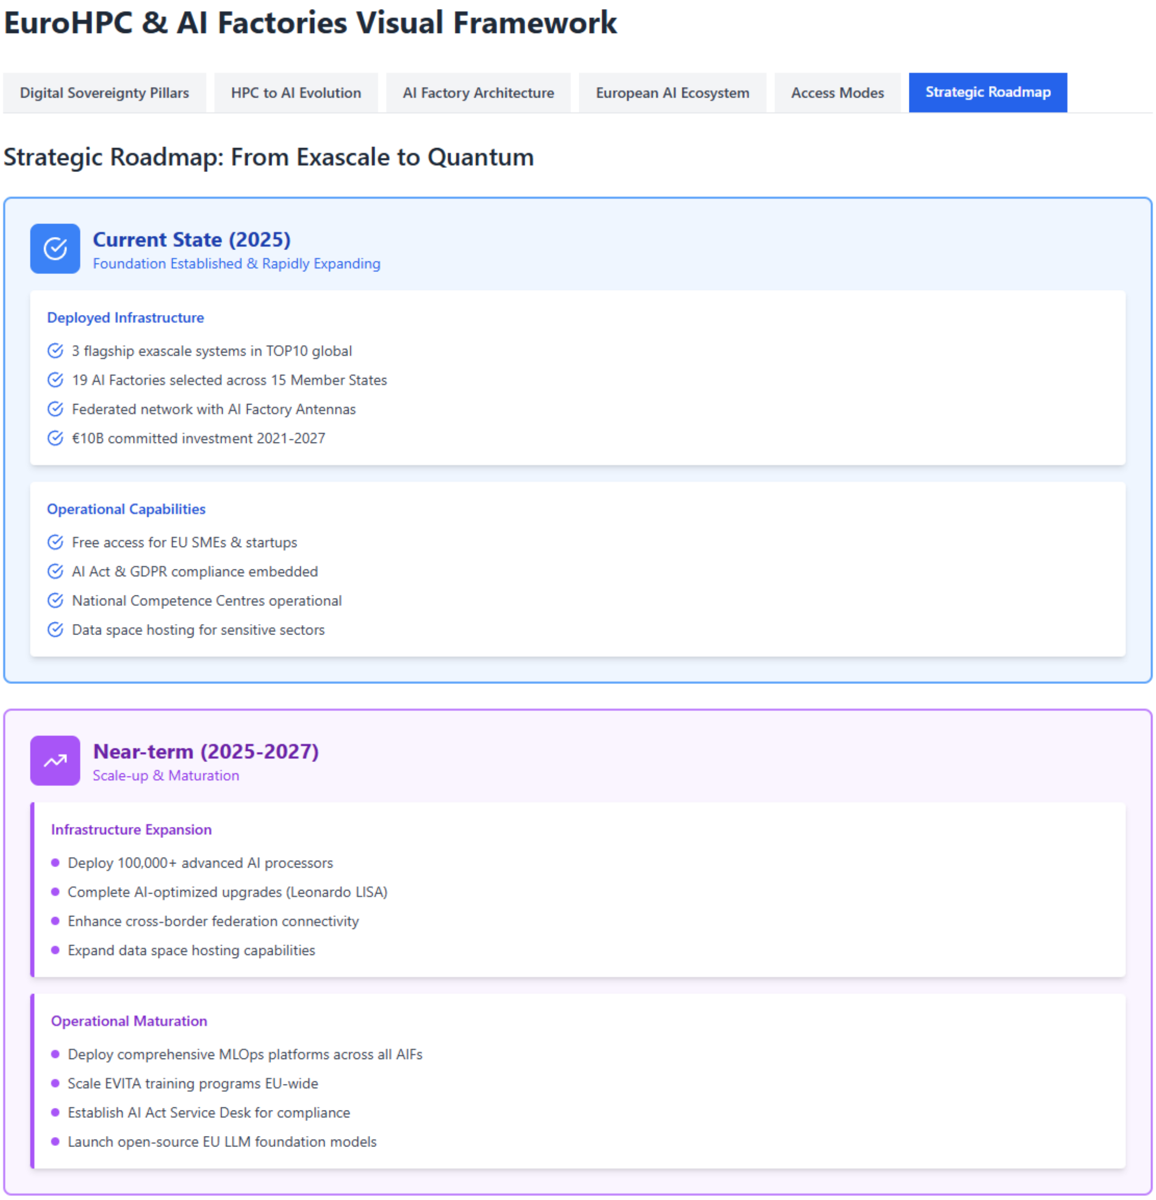

Furthermore, maintaining a competitive edge requires a nuanced approach to workload scaling. While the on-premises AI Factories prioritize sovereignty and compliance, future strategy involves sophisticated hybrid cloud integration. This will enable seamless workload distribution between the sovereign on-premises AIFs and commercial cloud resources, ensuring elasticity and responsiveness for periods of extreme or dynamic demand while strictly maintaining data sovereignty requirements.[6] The EuroHPC federated ecosystem is actively set up to foster close cooperation with platforms like GAIA-X, which focuses on industrial and commercial services, thus stimulating the utilization of JU infrastructures by the private sector, especially SMEs, through industry-friendly federation services.[18]

Medium-term (2027-2030)
Sovereignty Consolidation & Next-Gen Tech

Hardware Sovereignty

- **EPI Phase 2:** Deploy EU-designed processors in production
- Reduce dependency on non-EU silicon
- ARM/RISC-V accelerators for AI workloads
- Establish EU semiconductor supply chain

Quantum Integration

- Deploy industrial-scale quantum platforms
- Integrate quantum into federated HPC ecosystem
- Hybrid classical-quantum workflows
- Quantum-safe cryptographic implementations

Long-term (2030+)
Global Leadership & Complete Autonomy

Full Stack Sovereignty

✓ Indigenous silicon (EPI mature)
✓ EU-owned data spaces
✓ Open-source software stacks
✓ Zero external dependencies

AI Excellence

✓ World-leading trustworthy AI
✓ Globally competitive LLMs
✓ Human-centric AI standards
✓ Export European AI values

Hybrid Ecosystem

✓ Quantum-HPC-AI integration
✓ Edge-to-exascale continuum
✓ Dynamic workload distribution
✓ Sustainable operations

7. Conclusion: Securing Europe's Computational Future

The establishment and rapid expansion of the EuroHPC AI Factories represent a strategic inflection point for Europe, successfully pivoting the continent's immense exascale computational investment to meet the immediate and unique demands of generative AI development. This initiative is the most substantial and practical step taken to date toward mitigating critical computational dependencies that otherwise risk compromising Europe's technological autonomy.

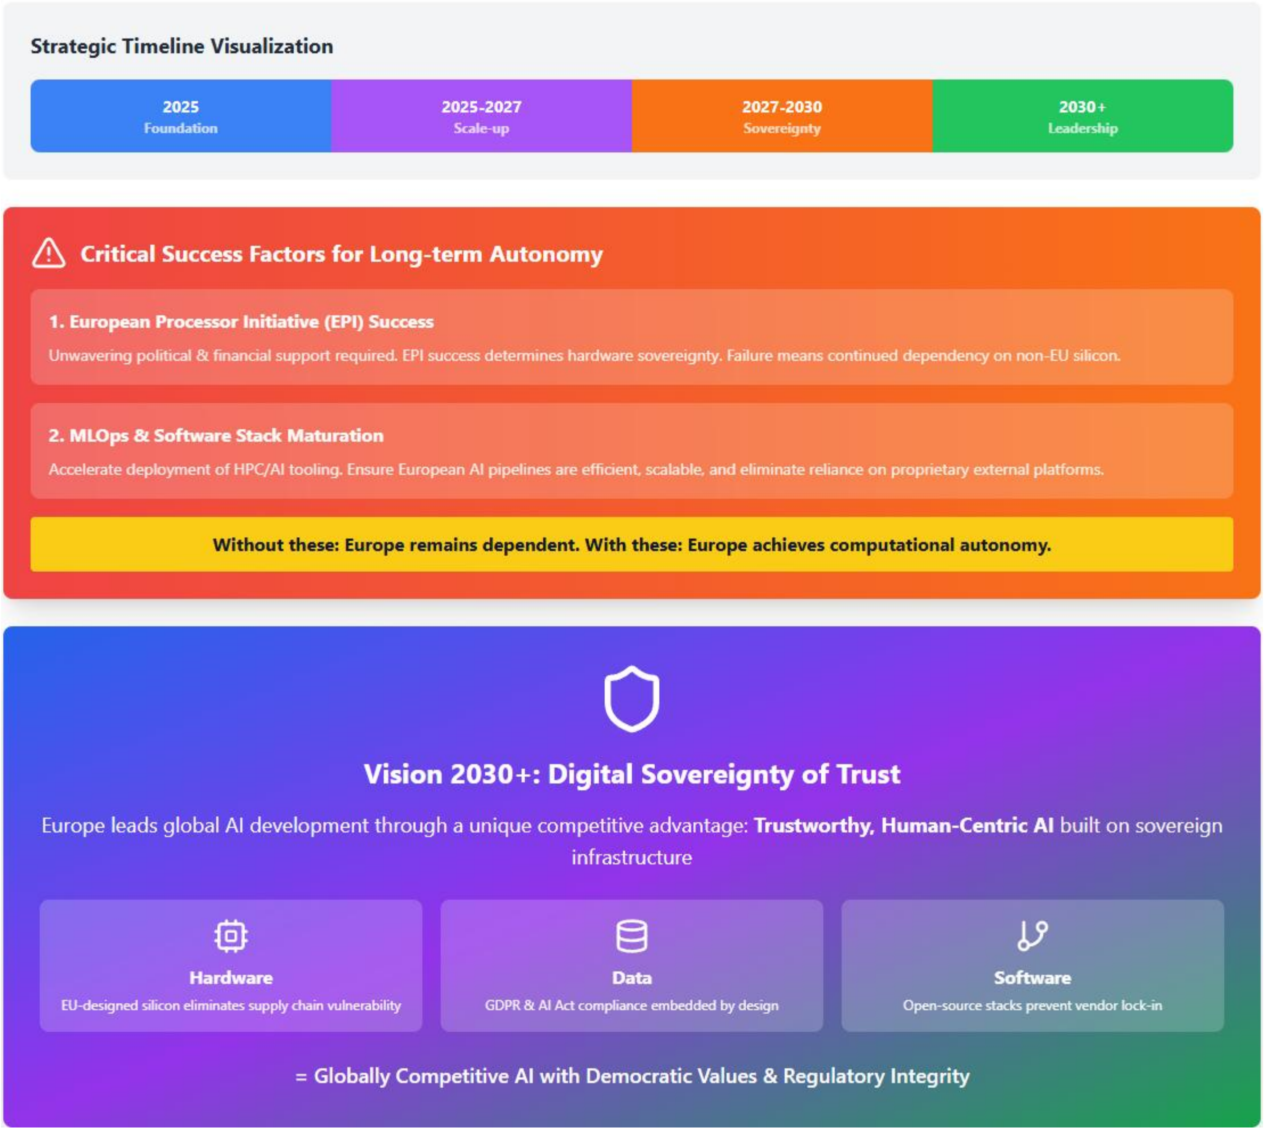

The core competitive advantage of this European approach lies in the enforced fusion of computational capacity with regulatory governance. By mandating the development of advanced AI models within an ecosystem that is structurally compliant with the AI Act and GDPR, Europe is actively building a **Digital Sovereignty of Trust**. This strategy ensures that European technological outputs are globally competitive not only in performance but also in their alignment with democratic values, privacy standards, and regulatory integrity.

To solidify comprehensive long-term autonomy, continued strategic commitment remains necessary in two key areas. First, unwavering political and financial support for the European Processor Initiative (EPI) is essential. The success of EPI determines whether the hardware component of sovereignty—the foundational silicon layer—can eventually be controlled domestically, moving beyond the current, calculated reliance on non-EU components. Second, the maturation and broad deployment of MLOps capabilities and HPC/AI tooling within the AI Factories must be accelerated. This operational maturity will ensure that European AI development pipelines are efficient, scalable, and resilient, eliminating reliance on proprietary external software stacks. By sustaining these coordinated efforts across hardware, data, software, and skills, the EuroHPC JU guarantees that the computational future of Europe is built on a foundation of control, resilience, and globally competitive power.

For more information about Scientific High Performance Computing and Digital Sovereignty contact our specialized HPC consulting team at ATOMIX

Quoted works and sources

1. Understanding European Tech Sovereignty: Challenges and Opportunities - Hivenet, https://www.hivenet.com/post/understanding-european-tech-sovereignty-challenges-and-opportunities
2. Discover EuroHPC JU - European Union, https://www.eurohpc-ju.europa.eu/about/discover-eurohpc-ju_en
3. EuroHPC JU - FCT, https://www.fct.pt/en/internacional/espaco-europeu-de-investigacao/eurohpc-ju/
4. European High Performance Computing Joint Undertaking - EuroHPC JU, https://digital-strategy.ec.europa.eu/en/policies/high-performance-computing-joint-undertaking
5. AI Factories | Shaping Europe's digital future - European Union, https://digital-strategy.ec.europa.eu/en/policies/ai-factories
6. AI Readiness Assessment.docx
7. EuroHPC JU Opens Quantum Grand Challenge Call - European Union, https://www.eurohpc-ju.europa.eu/eurohpc-ju-opens-quantum-grand-challenge-call-2025-10-14_en
8. The European Union's AI Factories, https://www.interface-eu.org/publications/ai-factories
9. Our Supercomputers - EuroHPC JU - European Union, https://www.eurohpc-ju.europa.eu/supercomputers/our-supercomputers_en
10. EuroHPC Supercomputers Put Europe at the Forefront of Global Supercomputing, https://www.eurohpc-ju.europa.eu/eurohpc-supercomputers-put-europe-forefront-global-supercomputing-2025-06-10_en
11. JAIF - JUPITER AI Factory - Forschungszentrum Jülich, https://www.fz-juelich.de/en/ias/jsc/jupiter/jaif-jupiter-ai-factory
12. EPI - European Processor Initiative, https://www.european-processor-initiative.eu/project/epi/
13. EPI - EuroHPC JU - European Union, https://www.eurohpc-ju.europa.eu/research-innovation/our-projects/epi_en

14. European AI Factories are a critical battlefield for tech sovereignty – Let's not hand an easy victory to Big Tech! | SovereignEdge.EU, https://sovereignedge.eu/blog/european-ai-factories-are-a-critical-battlefield-for-tech-sovereignty/
15. The EuroHPC JU Selects Six Additional AI Factories to Expand Europe's AI Capabilities, https://www.eurohpc-ju.europa.eu/eurohpc-ju-selects-six-additional-ai-factories-expand-europes-ai-capabilities-2025-10-10_en
16. EuroHPC JU Selects AI Factory Antennas to Broaden the AI Factories Initiative, https://www.eurohpc-ju.europa.eu/eurohpc-ju-selects-ai-factory-antennas-broaden-ai-factories-initiative-2025-10-13_en
17. TrustLLM - Large Language Model Technology for Europe - EuroHPC JU - European Union, https://www.eurohpc-ju.europa.eu/trustllm-large-language-model-technology-europe_en
18. EuroHPC JOINT UNDERTAKING DECISION OF THE GOVERNING BOARD OF THE EuroHPC JOINT UNDERTAKING No 8/2023 Approving the Multi-Annual, https://www.eurohpc-ju.europa.eu/document/download/ea5b58d8-0270-4247-ba5e-738173af4f1f_en?filename=Decision%2008_2023_%20Amendment%20MASP%202021-2027_0.pdf
19. REGULATIONS - EuroHPC JU - European Union, https://www.eurohpc-ju.europa.eu/system/files/2022-03/uriserv_OJ.L_.2021.256.01.0003.01.ENG_EN_TXT.pdf
20. EuroHPC Joint Undertaking Multi-Annual Strategic Programme (MASP 2021 -2027), https://www.eurohpc-ju.europa.eu/document/download/2a5db80d-6bd3-4ea9-9418-4b49cadcf040_en?filename=20240219_For_GB_Decision%20Amendment%20MASP%202021-2027%20v220324%20Final%201%20%281%29.pdf
21. Federation of Computing Infrastructures | ETP4HPC, http://etp4hpc.eu/wp-content/uploads/2025/10/ETP4HPC_Federation-of-computing-infrastructures_DRAFT-All-Rights-Reserved.pdf
22. Privacy Policy - EuroHPC JU - European Union, https://www.eurohpc-ju.europa.eu/privacy-policy_en
23. AI Act | Shaping Europe's digital future - European Union, https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai
24. Europe's AI Policy Momentum: From Simplification to Strategy and Infrastructure - MIAI, https://ai-regulation.com/europes-ai-policy-momentum-from-simplification-to-strategy-and-infrastructure/
25. EuroStack – A European Alternative for Digital Sovereignty - University College London, https://www.ucl.ac.uk/bartlett/sites/bartlett/files/eurostack.pdf
26. National Competence Centres for High Performance Computing - EuroHPC JU, https://www.eurohpc-ju.europa.eu/national-competence-centres-high-performance-computing-1_en
27. AI Factories Access Modes - EuroHPC JU - European Union, https://www.eurohpc-ju.europa.eu/ai-factories/ai-factories-access-modes_en
28. EuroHPC: Supercomputing access for AI development - Mobility Innovation Marketplace, https://marketplace.eiturbanmobility.eu/opportunities/eurohpc-supercomputing-access-for-ai-development
29. National Competence Centres | High-Performance Computing in Europe, https://hpc-portal.eu/nccs
30. The EuroHPC Virtual Training Academy (EVITA) Project: Developing coordinated HPC training across Europe with recognised qualifications, https://www.eurohpc-

ju.europa.eu/eurohpc-virtual-training-academy-evita-project-developing-coordinated-hpc-training-across-europe-2025-04-14_en
31. EuroHPC JU: Homepage, https://www.eurohpc-ju.europa.eu/index_en