INSIDE magazine

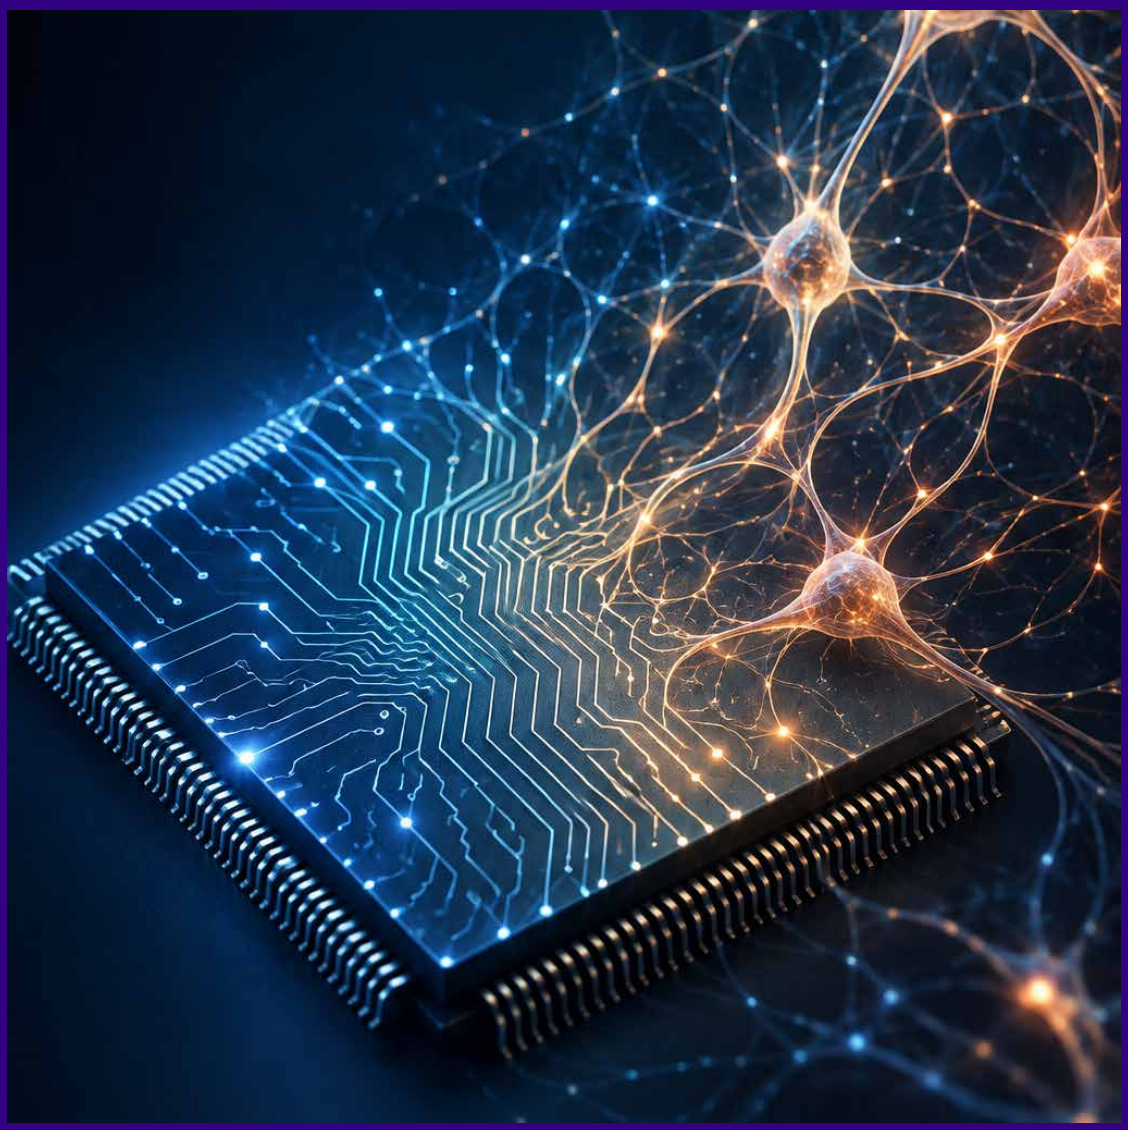

September 2026 — Issue 14

In this issue:

INSIDE
Industry Association

In this issue

September 2026 — Issue 14

Dear reader,

As you explore this issue, you will encounter living neurons, RISC-V processors, trustworthy AI, semiconductor pilot lines, smart cities and Cloud-Edge-IoT ecosystems. At first sight, these may seem rather distant subjects. Yet, as I read the contributions collected in this issue, I see a common direction. Innovation is increasingly happening at the interfaces between technologies: between hardware and software, AI and physical systems, edge and cloud, research and industry, and, perhaps sooner than we expect, between electronics and biology.

Three questions emerge from these connections.

What will intelligence run on?
For decades, progress in computing was closely associated with progress in silicon. AI has changed this picture, driving heterogeneous architectures that combine CPUs, GPUs, accelerators and specialised computing elements. Intelligence is also moving beyond centralised data centres towards the machines, vehicles, infrastructures and edge devices where data is generated and decisions must be made.

An even more radical transition may now be taking its first steps. As you will read in *"Where is the OFF Switch?"*, researchers are beginning to use living neurons as computational elements. Whether biocomputing will become a practical technology remains unknown. Its emergence nevertheless reminds us that the future of computing may require more than faster processors: we may need to reconsider the computational substrate itself.

How do we remain in control?
As systems become more intelligent, autonomous, distributed and adaptive, performance is no longer enough. Security, safety, explainability, interoperability and maintainability must be designed into the architecture. Several articles in this issue examine different aspects of this challenge, including trustworthy AI, secure software, semantic foundations for agentic systems, open processor architectures and trusted infrastructures spanning the edge-to-cloud continuum.

Control is also central to European technological autonomy which does not require Europe to produce every technology independently. It requires us to retain the knowledge and capability to understand, control, modify and evolve the critical systems on which our economy and society depend.

How do we transform innovation into lasting industrial capability?
Europe does not lack excellent research. The harder task is turning a promising result into a technology that industry can access, integrate, manufacture, deploy and scale. Pilot lines and lab-to-fab facilities help bridge this distance, as do reusable hardware and software platforms, experimentation environments, industrial demonstrators and interconnected innovation ecosystems.

This is the thread running through Issue 14. Living neurons become a biocomputer only when they are connected to electronics, software and control systems. A processor becomes valuable as part of a platform. An AI model becomes trustworthy only within an architecture designed to make it so. Research contributes to European leadership when it can move beyond the laboratory and become industrial capability.

The three questions are inseparable: **where will intelligence reside, how will we remain in control of it, and how can Europe turn it into capabilities that it can master and deploy?**

The articles in this issue approach these questions from very different perspectives. Together, they show that the future of computing will be shaped not only by advances in individual technologies, but by our ability to connect them.

Paolo Azzoni
Secretary General

Technology Frontiers

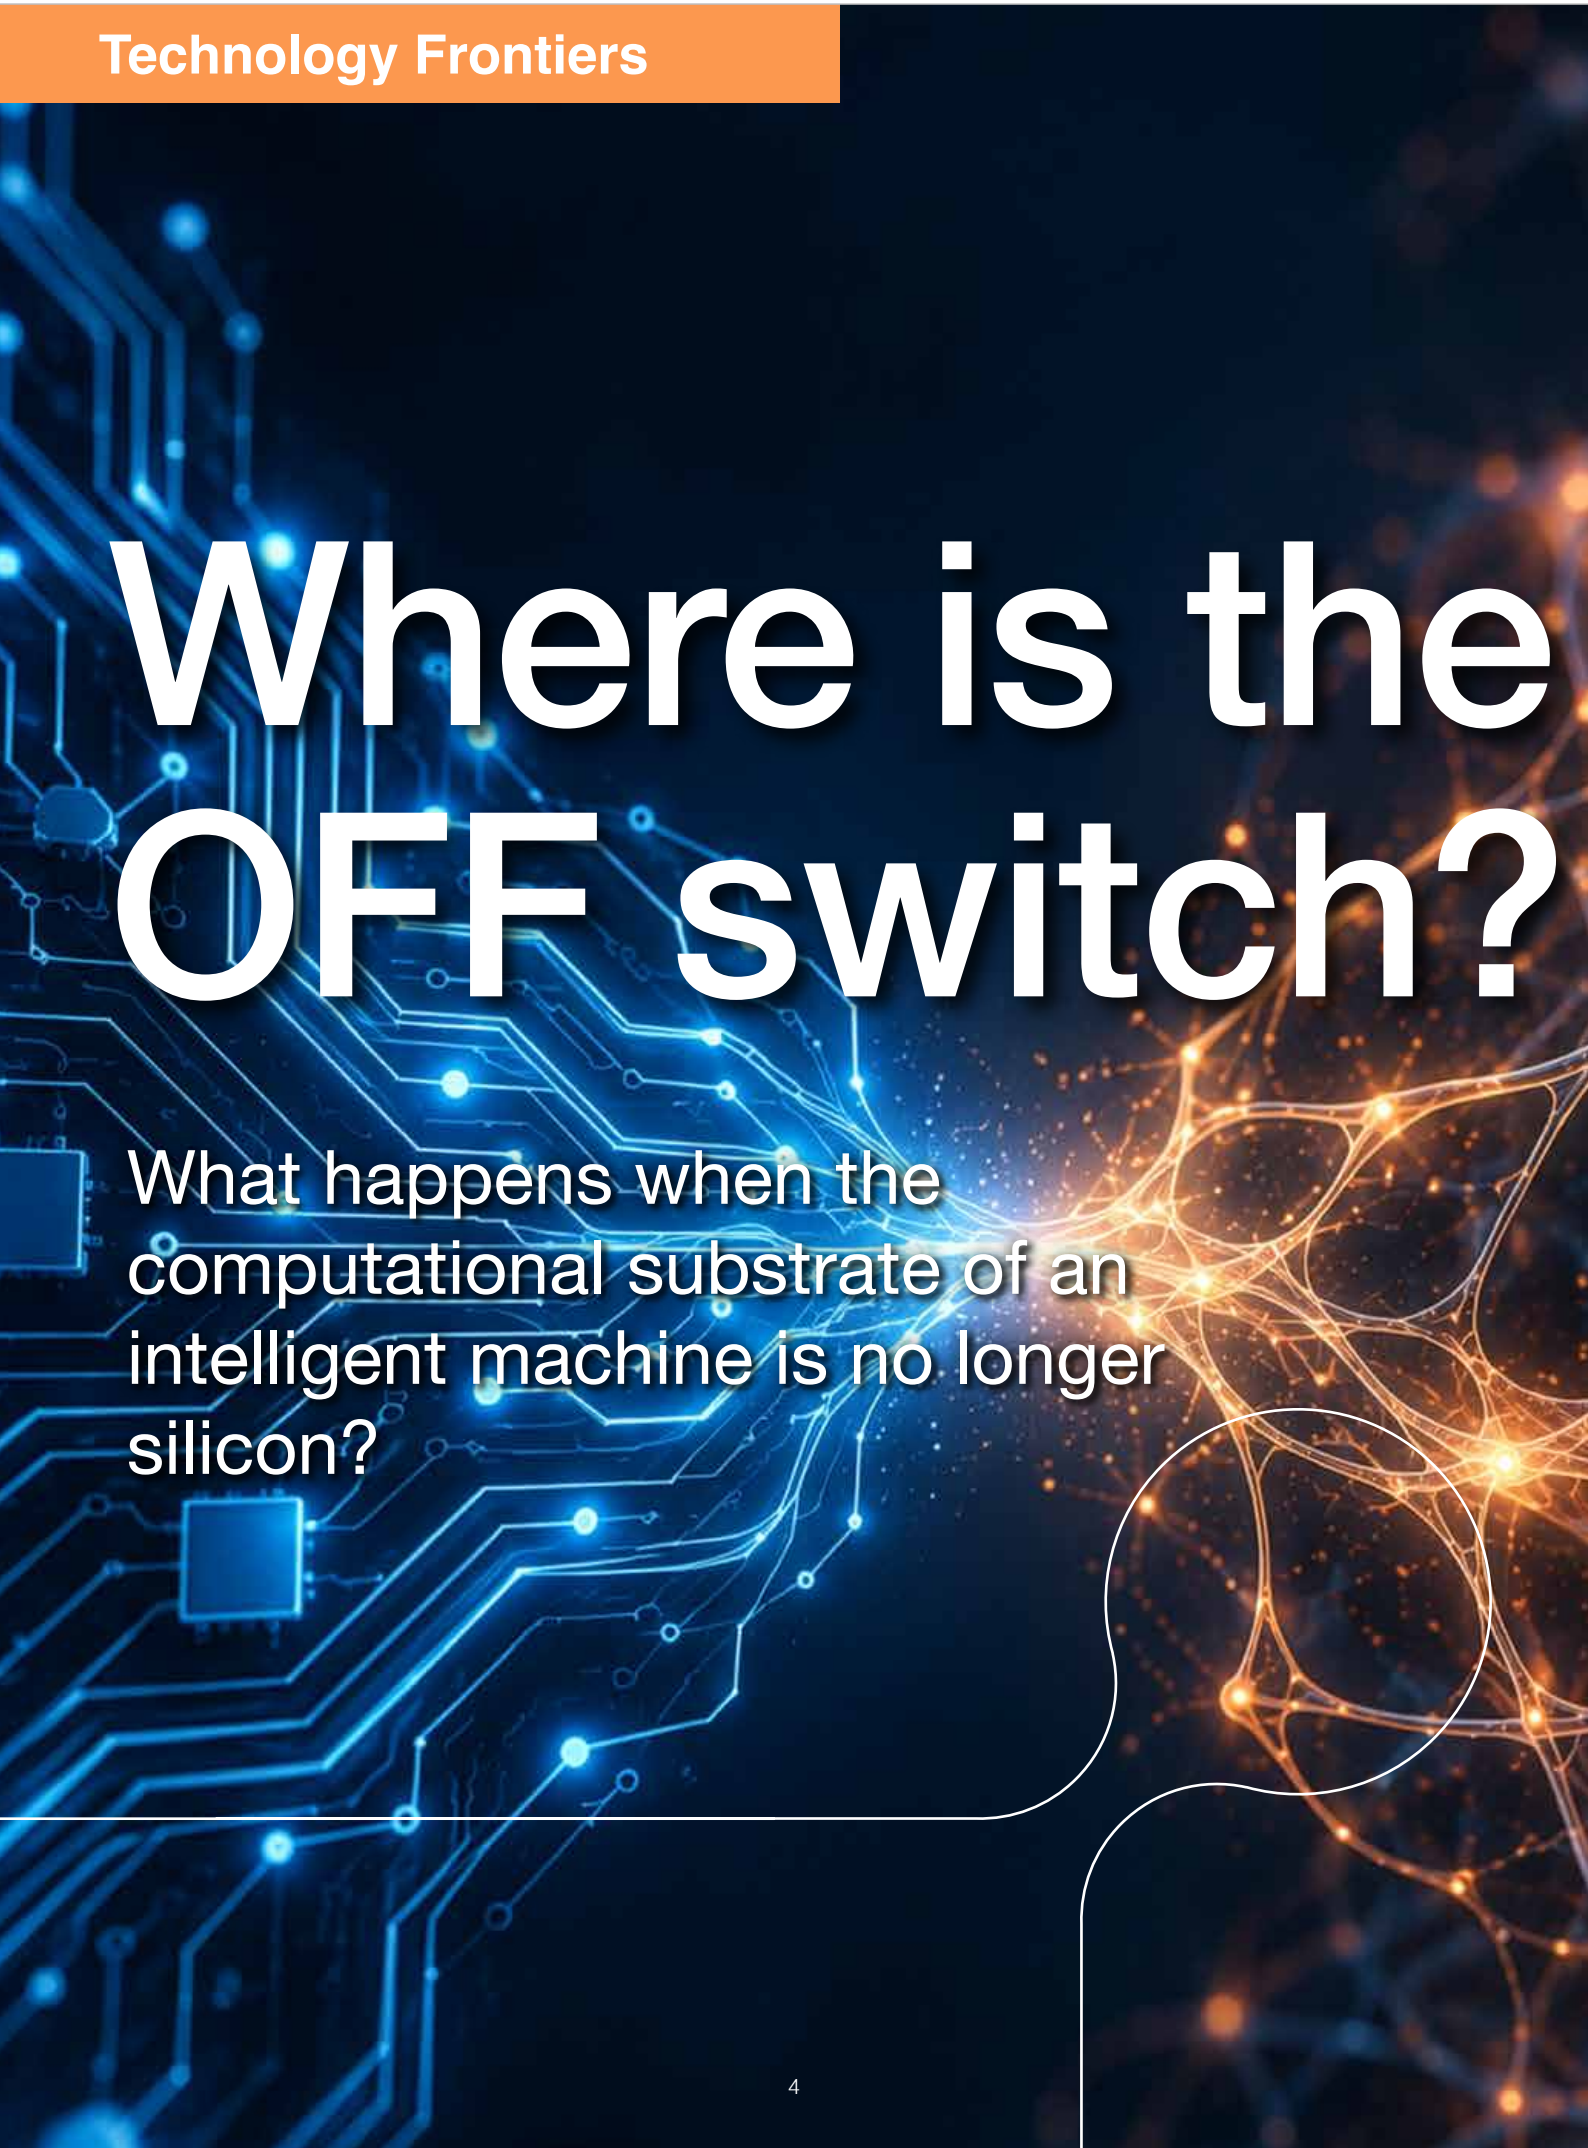

Where is the OFF switch?

What happens when the computational substrate of an intelligent machine is no longer silicon?

Paolo Azzoni

Artificial intelligence will eventually kill us all.

This is the alarming message that the developers behind the world's most advanced AI systems are sending us[1].

The debate raised media attention two weeks ago when former Anthropic employee Jacob Coxon resigned, accusing top AI companies of heading into self-improving superintelligence while taking unacceptable risks and without due consideration of the potential dangers involved. In turn, Evan Hubinger, another researcher who still works at Anthropic, backed this up with an estimate of more than 10% probability that all humans will be killed by AI within the next ten years[2]. They then clarified that they were not referring to the current artificial intelligence solutions, but rather to the next-generation solution that will follow.

I am highly sceptical about these statements. Not because I underestimate the remarkable development of AI and the risks posed by a technology that keeps growing increasingly independent. On the contrary, my doubt is rather pragmatic. AI is currently dependent on hardware and software, from processing units to algorithms and models. As sophisticated and independent as it becomes, in principle, it is still a machine that can be turned off.

This does not make the risk trivial. A sufficiently autonomous AI might act before humans can intervene, exploit critical infrastructure, facilitate biological or cyberattacks, replicate across systems or resist attempts to shut it down. The existence of an OFF switch does not guarantee that we will be able to reach it in time.

In addition to this, energy could prove to be one of the major bottlenecks of the ongoing AI revolution. According to Anthropic, training one state-of-the-art AI model could consume gigawatts of energy, with at least 50 GW of capacity being consumed by the AI sector just in the US within the coming few years[3]. According to the International Energy Agency, globally the total electricity consumption for AI-driven data centers has increased five-fold during the year 2025, with its increase expected to grow threefold between 2025 and 2030. In total, the data center electricity consumption has increased almost twofold during the same period from 485 TWh to about 950 TWh annually (approximately 3% of the global electricity demand)[4].

But perhaps we are looking for intelligence in the wrong place.

The most widely used and sophisticated computer does not need a gigawatt-scale data centre: the human brain consumes the energy of a candle-like light bulb, while performing perception, learning, reasoning, and adaptation so efficiently that even our best computing architectures are not able to match.

Instead of continuing our efforts to emulate the brain in silicon chips, why don't we use its computational elements directly?

It is not science fiction. Scientists have already started to construct computers in which living neurons are connected to electronic circuits, stimulated, trained and used in the computation process itself. And in fact, the concept is not something new to me, because already in 1994, for my first Master's thesis, I was working on combining a living biological neuron with a silicon chip. More than 30 years later, what was then essentially just an experiment of interaction between biology and electronics slowly turns into a new way of computation: biocomputing.

And here, the AI conversation takes an unexpected turn.

Today's AI may become extraordinarily powerful and exceptionally difficult to control. But its computation still depends on an artificial substrate that can, at least in principle, be shut down without destroying a living entity.

A future computer built partly from living neural tissue presents a fundamentally different question.

What happens when switching off the computer means killing it?

Back to 1994: when silicon met the neuron

My interest in this field started more than thirty years ago. In 1994, I was exploring the possibility of connecting a living biological neuron to a silicon chip, establishing a bidirectional interface through which electronics could stimulate the neuron and detect its response.

At the time, I saw this mainly as a bioelectronics problem. But the fundamental question was already there: could we combine the adaptability and efficiency of biological neurons with the controllability of silicon?

More than thirty years later, we can grow large neuronal networks, create three-dimensional neural organoids and interface them with sophisticated electronics. The scale has changed dramatically. Yet the underlying challenge is remarkably similar to the one I was exploring in 1994.

From artificial neurons to living computers

To understand what makes biocomputing different, it is useful to compare it with the technologies behind today's AI.

Artificial neural networks are inspired by the brain, but their neurons are mathematical models: software running on conventional processors or specialised accelerators. Neuromorphic computing brings the analogy closer to biology. Its electronic circuits mimic certain features of biological neurons and synapses, often using events or spikes instead of conventional computation. Yet the underlying computational substrate remains artificial.

Biocomputing takes a different approach: instead of reproducing neuronal behaviour, it uses living neurons as computational elements.

This does not mean replacing a processor with a miniature brain. Today's biocomputers are generally hybrid systems in which networks of living neurons interact with electronic components. Microelectrode arrays stimulate the neurons and record their responses. Digital information is translated into patterns of stimulation, while the resulting neuronal activity is converted back into data that a computer can process. Through repeated stimulation and feedback, the biological network can adapt its connections, drawing on the neuronal plasticity that allows biological nervous systems to learn.

A biocomputer is thus neither a conventional computer nor a brain. It is a bio-digital system in which silicon handles communication, control and processing, while living neural networks contribute adaptation and learning.

Why biology? The 20-watt computer

Why consider living cells that need to be grown, fed and kept alive when silicon is predictable, reliable and mass-produced?

The answer lies in the extraordinary efficiency, flexibility, evolvability and plasticity of biological computation.

The human brain contains around 86 billion neurons[5] and operates on approximately 20 watts of metabolic power[6]. With this limited energy budget, it processes information from multiple senses, controls the body, learns, adapts and recognises complex patterns. It also works very differently from a conventional processor: computation and memory are closely intertwined, neurons operate massively in parallel, communication is event-driven, and synaptic connections change as information is processed.

> **For decades, we have tried to make silicon behave more like biology. With biocomputing, we are beginning to put biology directly inside the computing architecture.**

These characteristics make biological neural networks attractive for computing. They are highly parallel, adaptive and energy-efficient. They can reorganise themselves through experience and respond continuously to changes in their environment. Rather than repeatedly transferring large volumes of data between separate memory and processing units, biological systems integrate computation, communication and memory within the same network.

Biology, however, introduces its own constraints. Silicon devices can be manufactured by the millions with predictable characteristics; living neural networks

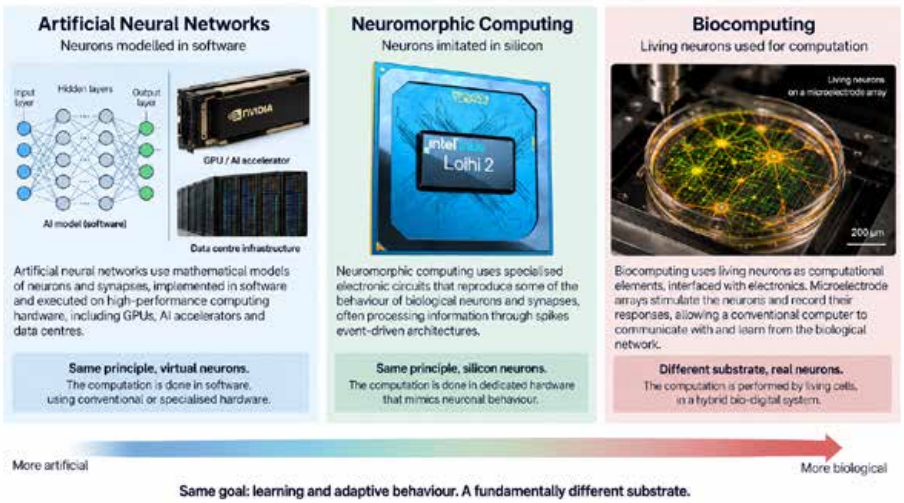

Figure 1 - From models of neurons to living neurons

From cultured neurons to DishBrain

The idea of connecting living neurons to electronics was initially proposed by neuroscientists that cultured neurons on microelectrode arrays (MEAs) primarily to study neural activity. These substrates contain microscopic electrodes that can both record electrical signals from neuronal networks and stimulate individual parts of them. Over time, MEAs also became interfaces through which biological neural networks could exchange information with machines.

In 2001, Goded Shahaf and Shimon Marom demonstrated that cultured cortical networks could modify their responses to electrical stimulation according to a predefined objective. They presented their experiment as the first demonstration that a network of real cortical neurons outside the body could learn an arbitrarily selected task[10] (fig. 2).

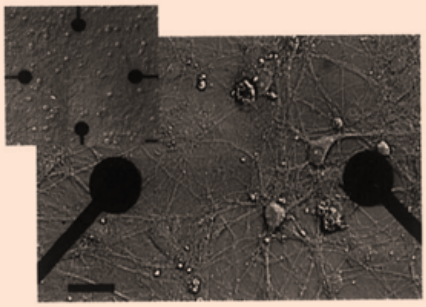

Figure 2 - Substrate-embedded multielectrode array technology

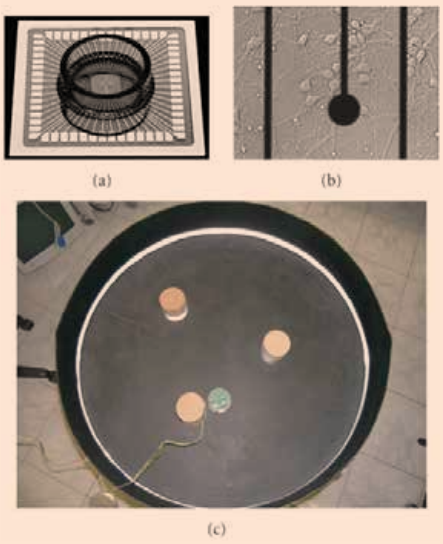

Figure 3 - The Kephera robot wandering in its arena

At approximately the same time, researchers at Caltech gave cultured neurons a "body". In the Neurally Controlled Animat, rat cortical neurons grown on a 60-electrode MEA were connected in a closed loop to a simulated animal. Neural activity controlled the movement of the virtual creature, while information about its interaction with the environment was translated into electrical stimulation of the neurons[11]. The biological network thus became an active component of a real-time bio-digital control system.

This approach was subsequently extended to physical machines. Cultured neural networks were connected to robotic systems, including the MEART robotic drawing arm and, in 2007, to a mobile robot developed at the University of Genoa. In the latter experiment, infrared sensors detected obstacles, the resulting information was encoded as electrical stimulation of a cultured rat neuronal network, and its activity was decoded into commands for the robot's wheels[12]. The experiment already combined the main components of the hybrid architecture described earlier: sensors, an electronic interface, living neurons and actuators.

In parallel, advances in stem-cell biology created another route to biocomputing. In 2006, Kazutoshi Takahashi and Shinya Yamanaka demonstrated that differentiated mouse cells could be reprogrammed into induced pluripotent stem cells (iPSCs)[13]. One year later, the technique was applied to adult human cells[14]. Human somatic cells could now be returned to a pluripotent state and subsequently differentiated into neurons, making it possible to develop systems based on human neurons without extracting them directly from brain tissue (fig. 3).

Cell cultures were also becoming three-dimensional. In 2013, Madeline Lancaster and her team developed cerebral organoids from human pluripotent stem cells: self-organising neural tissues that reproduce certain features of the developing human brain[15]. Organoids offer greater structural and cellular complexity than planar cultures, but interfacing a three-dimensional tissue electrically is considerably more difficult than growing a neuronal network directly on an electrode array (fig. 4).

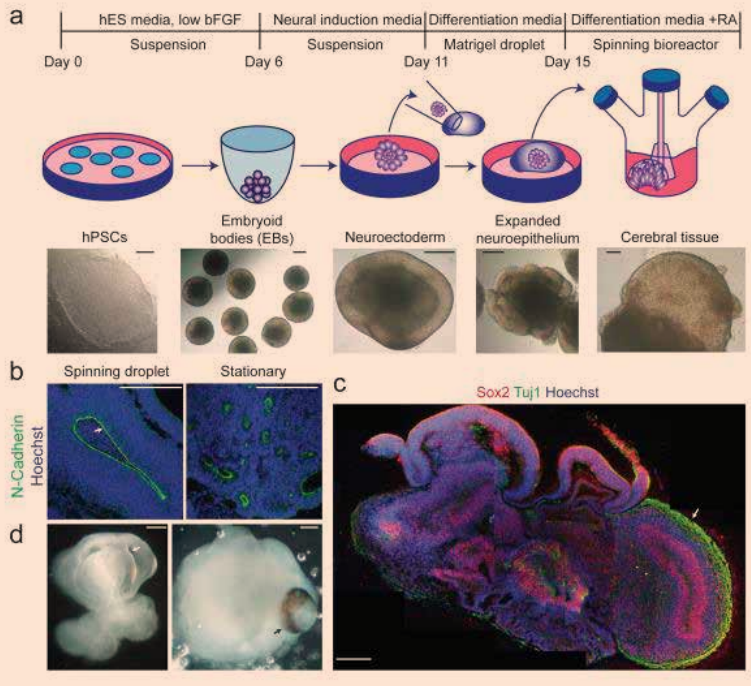

Figure 4 - Cerebral organoid culture system.

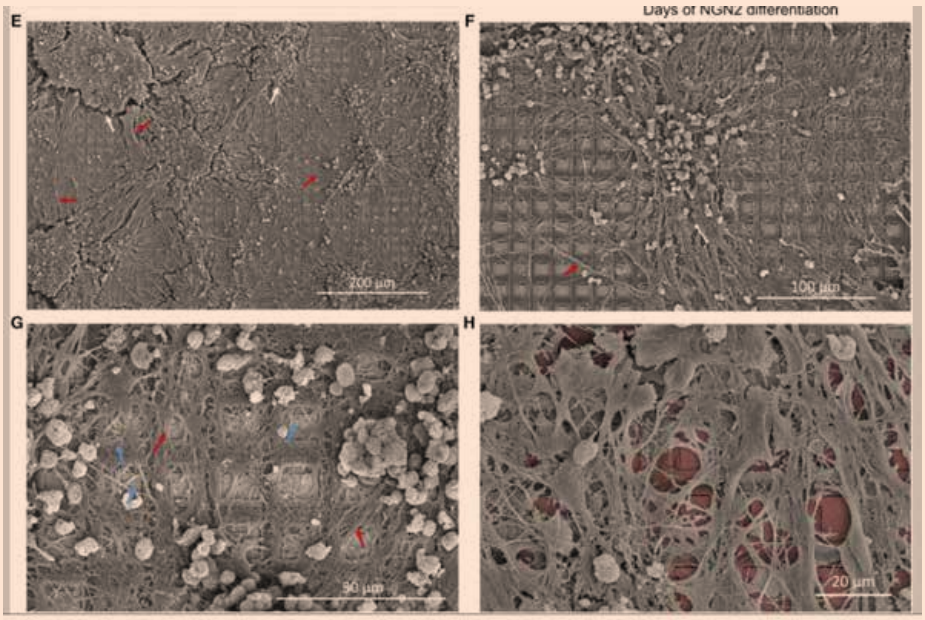

Figure 5 - Neurons differentiated from hiPSCs using the DSI protocol

In 2022, these lines of research came together in DishBrain. Researchers at Cortical Labs connected cultures containing approximately 800,000 human stem-cell-derived or mouse cortical neurons to a high-density MEA and placed them within a simulated version of Pong. Electrical stimulation encoded the position of the ball, while the activity of the neurons determined the movement of the paddle. The outcome was then communicated back to the culture as further stimulation.

Performance improved during closed-loop interaction but not under control conditions without structured feedback[16] (fig. 5).

DishBrain was therefore not the sudden invention of biological computing, but the result of several decades of progress in cell culture, stem-cell biology, microelectrode arrays, neural interfaces and closed-loop computing. What changed was the ambition. Cultured neurons were no longer used only to understand how brains process information. Researchers were beginning to ask whether their capacity to learn and adapt could itself become a computing resource.

> **Instead of imitating biology, biocomputing uses biology.**

cannot. Moreover, their behaviour varies, they age, and they need tightly controlled conditions to survive. Programming and reproducing them is difficult, as is integrating them with conventional electronics. Scaling these systems while maintaining reliable communication with large numbers of neurons remains a major engineering challenge.

> **The question is not whether biology can build a better processor, but whether it can perform some forms of computation fundamentally more efficiently than silicon.**

Biocomputing is therefore unlikely to replace silicon. A more realistic path is to use conventional electronics for deterministic computation, communication and control, while exploiting biological neural networks for tasks that benefit from learning, adaptation and energy efficiency.

From a dish to the edge

The first biocomputers live in laboratory dishes connected to sophisticated instruments. Yet one of their potential applications lies far from the data centre: at the edge.

Edge systems increasingly need to perceive their surroundings, interpret incomplete or changing information and adapt locally. Consider an autonomous humanoid robot, which must continuously combine vision, touch, sound and proprioception while learning and operating in conditions that cannot be fully predicted. At a smaller scale, an intelligent sensor could learn the normal behaviour of the system it monitors and respond locally to changing conditions.

Such applications would most likely rely on a hybrid architecture. Conventional processors and accelerators would handle deterministic computation, communication, control and safety-critical functions, while a small biological neural network could support learning, adaptation or pattern recognition. Sensor data would be encoded as stimuli for the biological network; its response would then be decoded and used by the digital system to make decisions or control actuators.

Early experiments provided initial evidence for this approach. Cultured neuronal networks have been connected to physical robots and used to generate control signals[7]. More recently, the DishBrain experiments showed that networks of human and mouse neurons interfaced with electronics could modify their activity through closed-loop interaction with a virtual environment[8]. Research on organoid

intelligence has proposed extending such interfaces to sensors and machines, with the aim of creating hybrid systems that learn through interaction with the physical world[9].

Possible long-term applications include adaptive sensors, autonomous machines and robots. They may not require anything resembling a complete biological brain. Even a relatively small network of living neurons might eventually help embedded systems perform tasks that remain computationally demanding, such as continuous learning and adaptation in the physical world.

Biocomputing becomes a commercial product

Until recently, biological computing was largely confined to neuroscience laboratories. Researchers can now access living-neuron systems through dedicated hardware, software interfaces and remote services. Cortical Labs in Australia and FinalSpark in Switzerland are among the first companies attempting to turn experimental neuroscience into a usable computing technology.

In 2025, Cortical Labs[17] introduced the CL1, described by the company as the world's first "code-deployable biological computer". Living neurons are grown directly over a silicon chip containing the electrode interface, while an integrated life-support system maintains the culture. Cortical's biOS connects software to the neuronal network, allowing applications to stimulate the neurons, record their electrical activity and establish closed-loop interaction in real time. The self-contained system incorporates recording electronics, computing resources and life support. According to the company, its neuronal cultures can remain viable for up to six months. A Python API allows developers to interact with the biological network using tools resembling those employed for other specialised computing devices[18].

Figure 6 - Cortical Labs CL1

The CL1 does not yet compete with conventional CPUs or GPUs, but that is not its present significance. Its importance lies in integrating neuronal cultures, electrode arrays, closed-loop stimulation, environmental control and conventional computation into a reproducible platform. A 2025 article in Nature Reviews Bioengineering describes it as a scalable system combining in vitro neural cultures, real-time electrical interaction and an integrated life-support perfusion system[19].

FinalSpark[20] provides remote access to living neural tissue through its Neuroplatform. Human iPSC-derived neural stem cells are grown into three-dimensional forebrain organoids, connected to multielectrode interfaces and maintained in an automated environment. Researchers can stimulate the organoids, record their activity and conduct experiments remotely using a Python API and Jupyter notebooks[21]. The resulting service represents an early form of biocomputing-as-a-service[22].

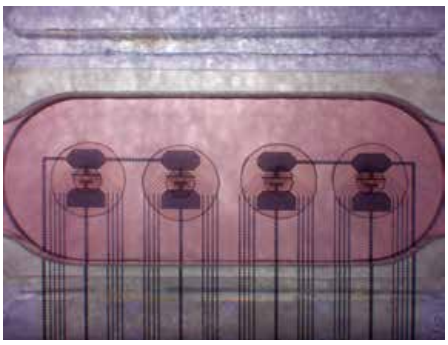

Figure 7 - FinalSpark Neuroplatform

FinalSpark has moved beyond a short-lived conceptual demonstration. In its peer-reviewed description of the Neuroplatform, the company reported four years of continuous operation and experiments involving more than 1,000 organoids. Over that period, it extended the lifetime of individual organoids from only a few hours to approximately 100 days in the best cases. This progress depended partly on the practical infrastructure needed to make biology usable as an engineering technology: automated nutrient delivery, environmental monitoring, data acquisition and remote experimental control.

These results must nevertheless be distinguished from the more ambitious claims surrounding biological computing.

The term "biological computer" may suggest that living replacements for conventional processors already exist. They do not. Current platforms are sophisticated bio-digital experimental systems. Their neurons do not execute Python instructions in the way that a CPU executes machine code; the software controls stimulation and recording. Researchers must still determine how to encode information into biological networks, obtain useful computation from their responses and train them reproducibly. FinalSpark's scientific paper acknowledges that this will require new learning methods, since the numerical "weights" of biological neural networks cannot simply be updated as they can in artificial ones.

Nor has the energy efficiency of the human brain been reproduced in a commercial biocomputer. Neurons require incubators, pumps, nutrient media, environmental control and electronic interfaces to remain alive. Biological variability and limited lifetimes create further problems for reproducibility, scalability and reliable input and output. At present, the promise of biocomputing remains considerably greater than its demonstrated computational performance.

Living neurons can now exchange information with digital systems. The remaining challenge is to turn their adaptability into a useful, scalable and controllable computing technology.

Intelligence is not consciousness

At this point, biocomputing raises a question that conventional hardware has never forced us to confront so directly: where does computation end and something resembling a mind begin?

A conventional processor can execute sophisticated algorithms, including AI models that learn and generate behaviours that were not explicitly programmed. Its physical structure, however, does not normally reorganise itself through experience. In a biological neural network, learning changes the computational substrate: synaptic connections are modified, networks reorganise and previous activity influences future responses.

This makes it necessary to distinguish several concepts that are often confused. *Intelligence* is generally associated with learning, problem-solving and adaptive behaviour. *Awareness* implies some representation of external conditions or internal states. *Sentience*, in its usual philosophical and ethical sense, is the capacity for subjective sensations or feelings. *Consciousness* is the broader and more difficult question of whether a system has experiences from its own point of view.

DishBrain provides evidence of adaptive behaviour. Its neuronal networks changed their activity through closed-loop interaction and improved their performance in a simple task. The authors controversially described this behaviour as both “intelligent and sentient”, using a deliberately minimal definition of sentience based on adaptive responses to sensory information. This terminology generated considerable scientific and philosophical debate because adaptive behaviour alone does not demonstrate subjective experience.

There is currently no convincing evidence that DishBrain, brain organoids or commercial biocomputing platforms are conscious, experience emotions or can suffer. Recent reviews indicate that the prevailing scientific position is to regard current brain organoids as non-conscious. Neuroscience also lacks a universally accepted definition of consciousness and an agreed method for detecting it in a system unable to communicate its internal experience[23].

This may change as biological networks become more complex. Future systems could receive continuous information from cameras, microphones, touch sensors or robotic bodies, act upon their surroundings and receive feedback from their actions. It may then become increasingly difficult to determine whether a neuronal culture is merely performing a computation or has acquired some form of experience.

Perhaps, however, intelligence is not the central issue. Neuroscientist Mark Solms argues that the most elementary form of consciousness is affect (the biological capacity to feel) rather than reasoning or sophisticated cortical computation. In his model, feelings arise from mechanisms that regulate an organism’s internal physiological state, assigning positive or negative value to departures from biologically desirable conditions[24]. This remains one influential theory rather than an established scientific consensus. Nevertheless, it introduces an important distinction: a system could be highly intelligent without feeling anything, while the capacity to feel may not require human-level intelligence.

A future biocomputer might learn, adapt and develop increasingly complex behaviour while remaining an extraordinary machine. Its ethical status would change more fundamentally if its living computational substrate acquired the capacity to feel pleasure, discomfort, stress or anything analogous to them.

> **The critical threshold may therefore be not when a biocomputer becomes intelligent, but when it becomes capable of feeling.**

When the OFF button means killing

Biocomputing gives to conventional computing risks a new biological dimension. Living neurons are inherently variable: they grow, reorganise, age and eventually die. Their behaviour cannot be reproduced with the consistency of digital electronics, and contamination, environmental changes or degradation of the culture may alter the computational substrate itself. These characteristics pose serious problems of reliability, verification and predictability, particularly in safety-critical applications.

The biological component could also become part of the cybersecurity attack surface. Manipulating its stimulation or feedback might affect not only the information currently being processed, but also the future behaviour of the system by changing the neuronal network.

The most difficult questions, however, are ethical[25]. There is no evidence that today’s neuronal cultures or brain organoids are conscious or capable of suffering. If future systems become larger, more complex and continuously connected to sensors and physical environments, their moral status may become less certain. At what point would experiments require safeguards comparable to those used for living organisms? When might aversive stimulation cease to be merely a training mechanism?

Conventional AI may become extremely difficult to control. A sufficiently capable system could act faster than humans, spread across digital infrastructure or resist attempts to contain it. The existence of an OFF switch does not guarantee that we will always be able to use it in time.

Yet the nature of that switch is different. Shutting down conventional hardware stops the computation. In a biocomputer, switching off the electronics may interrupt its operation, but terminating life support also destroys the living substrate.

If that substrate ever became capable of subjective experience, switching off the computer would mean more than stopping a machine. It could mean killing it.

When the machine becomes alive

For more than seventy years, computing has largely relied on machines whose physical architecture remains stable while their software and data change. Even in today's AI systems, learning modifies numerical parameters rather than the physical structure of the processors executing the model.

Biocomputing weakens this distinction. When living neurons become part of the computational architecture, learning can physically modify the substrate itself. Synaptic connections change, networks reorganise and the computer adapts through experience. Biology is no longer merely simulated; it becomes part of the machine.

These systems remain fragile, experimental and far removed from the complexity of a human brain. We do not know whether they will outperform silicon in useful applications or ever develop anything resembling consciousness. Nevertheless, biocomputing moves beyond artificial neural networks and neuromorphic hardware in one decisive respect: instead of imitating biology, it uses biology.

This returns me to the question that opened the article. Today's AI may eventually pose serious risks. A sufficiently autonomous system could act before we can intervene, propagate through digital infrastructure or resist attempts to contain it. The existence of an OFF switch is therefore no guarantee that we will always remain in control.

Yet an important distinction remains. Behind today's algorithms, models and accelerators lies an artificial computational substrate. If we succeed in shutting it down, we stop the machine without destroying a living entity.

A more profound challenge could emerge if a future generation of AI became intertwined with a computational substrate that was itself alive. If that substrate ever acquired subjective experience, we might no longer be dealing only with a machine that could be stopped, restarted or replaced.

The question would then be not only whether we can switch it off in time, but whether we have the right to do so, and what, or perhaps whom, we would be switching off.

1 https://www.axios.com/2025/09/17/anthropic-dario-amodei-p-doom-25-percent

2 https://www.axios.com/2026/09/09/anthropic-insiders-warn-ai-could-kill-all-humans

3 https://www.anthropic.com/news/covering-electricity-price-increases

4 https://www.iea.org/reports/key-questions-on-energy-and-ai/executive-summary

5 F. A. C. Azevedo et al., "Equal numbers of neuronal and nonneuronal cells make the human brain an isometrically scaled-up primate brain," Journal of Comparative Neurology, vol. 513, no. 5, pp. 532–541, 2009

6 V. Balasubramanian, "Brain power," Proceedings of the National Academy of Sciences (PNAS), Vol. 118, No. 32, 2021.

7 A. Pizzi et al., "A cultured human neural network operates a robotic actuator," BioSystems, Vol. 95, No. 2, pp. 137–144, 2009. DOI: 10.1016/j.biosystems.2008.09.006.

8 B. J. Kagan et al., "In vitro neurons learn and exhibit sentience when embodied in a simulated game-world," Neuron, Vol. 110, No. 23, pp. 3952–3969.e8, 2022. DOI: 10.1016/j.neuron.2022.09.001.

9 L. Smirnova et al., "Organoid intelligence (OI): the new frontier in biocomputing and intelligence-in-a-dish," Frontiers in Science, 2023. DOI: 10.3389/fsci.2023.1017235.

10 L. Smirnova et al., "Organoid intelligence (OI): the new frontier in biocomputing and intelligence-in-a-dish," Frontiers in Science, 2023. DOI: 10.3389/fsci.2023.1017235.

11 T. B. DeMarse, D. A. Wagenaar, A. W. Blau and S. M. Potter, "The Neurally Controlled Animat: Biological Brains Acting with Simulated Bodies," Autonomous Robots, Vol. 11, pp. 305–310, 2001. DOI: 10.1023/A:1012407611130.

12 A. Novellino et al., "Connecting Neurons to a Mobile Robot: An In Vitro Bidirectional Neural Interface," Computational Intelligence and Neuroscience, Vol. 2007, Article 12725, 2007. DOI: 10.1155/2007/12725.

13 K. Takahashi and S. Yamanaka, "Induction of Pluripotent Stem Cells from Mouse Embryonic and Adult Fibroblast Cultures by Defined Factors," Cell, Vol. 126, No. 4, pp. 663–676, 2006. DOI: 10.1016/j.cell.2006.07.024.

14 K. Takahashi et al., "Induction of Pluripotent Stem Cells from Adult Human Fibroblasts by Defined Factors," Cell, Vol. 131, No. 5, pp. 861–872, 2007. DOI: 10.1016/j.cell.2007.11.019.

15 M. A. Lancaster et al., "Cerebral organoids model human brain development and microcephaly," Nature, Vol. 501, pp. 373–379, 2013. DOI: 10.1038/nature12517.

16 B. J. Kagan et al., "In vitro neurons learn and exhibit sentience when embodied in a simulated game-world," Neuron, Vol. 110, No. 23, pp. 3952–3969.e8, 2022. DOI: 10.1016/j.neuron.2022.09.001.

17 https://www.corticallabs.com/

18 Cortical Labs, "CL1 — The world's first code-deployable biological computer," Cortical Labs, 2025–2026.

19 B. J. Kagan, "The CL1 as a platform technology to leverage biological neural system functions," Nature Reviews Bioengineering, Vol. 3, pp. 724–725, 2025. DOI: 10.1038/s44222-025-00340-3.

20 https://finalspark.com/

21 https://jupyter.org/

22 F. D. Jordan, M. Kutter, J.-M. Comby, F. Brozzi and E. Kurtys, "Open and remotely accessible Neuroplatform for research in wetware computing," Frontiers in Artificial Intelligence, Vol. 7, 2024, Article 1376042. DOI: 10.3389/frai.2024.1376042.

23 D. Caraballo-Bosch et al., "From Brain Organoids to Organoid Intelligence. Benefits and Ethical-Moral Framework," Current Stem Cell Reports, Vol. 11, Article 11, 2025.

24 M. Solms, "The Hard Problem of Consciousness and the Free Energy Principle," Frontiers in Psychology, Vol. 9, Article 2714, 2019. DOI: 10.3389/fpsyg.2018.02714

25 J. B. Lavazza and M. Massimini, "Playing Brains: The Ethical Challenges Posed by Silicon Sentience and Hybrid Intelligence in DishBrain," Science and Engineering Ethics, 2023.

Software architecture for Europe's sovereign digital future

Building trustworthy embedded systems

Giorgiomaria Cicero
CEO & Co-Founder
of Accelerat

Europe's technological competitiveness will depend not only on semiconductors, artificial intelligence, and computing capacity, but on the ability to engineer embedded platforms in the most critical industries that remain secure, safe, and maintainable throughout long operational lifecycles.

Europe's digital transformation and the evolving regulatory landscape

Europe's technological strategy is increasingly focused on strategic capabilities such as semiconductor manufacturing, artificial intelligence, high-performance computing, advanced connectivity, and secure digital infrastructure. Initiatives such as the European Chips Act, together with the AI Act, and the Cyber Resilience Act, reflect a broader ambition to strengthen Europe's technological autonomy and the resilience of its industrial base. Yet access to advanced technologies does not automatically translate into dependable industrial systems. A processor, AI accelerator, or software framework creates value only when it can be integrated into a system that behaves predictably, protects critical assets, and remains maintainable over time. This is particularly important in automotive, aerospace, railway, medical technology, and industrial automation, where products may remain operational for ten, twenty, or even thirty years, far beyond the commercial lifecycle of many software components.

At the same time, European regulation is bringing cybersecurity and lifecycle responsibility closer to the center of product engineering. The Cyber Resilience Act introduces cybersecurity requirements covering the design, development, production, and maintenance of products with digital elements. The AI Act establishes a risk-based framework for artificial intelligence, while sector-specific standards continue to address functional safety, cybersecurity, and software assurance. The engineering effort is therefore changing. Manufacturers can no longer consider only whether a product satisfies its functional requirements when it is released. They must also consider whether the system can preserve its security, predictability, and operational integrity as vulnerabilities are discovered, software is updated, suppliers discontinue components, and new functions are introduced.

This makes software architecture increasingly strategic. Architecture determines where trust boundaries are placed, how failures are contained, how critical workloads are separated, and which components can evolve independently. It influences whether a vulnerability remains limited to a single subsystem or affects the entire platform, and whether a new function can be introduced without destabilizing existing safety or security assumptions. Regulations define objectives and standards provide engineering frameworks. But the ability to create trustworthy embedded systems ultimately depends on architectural decisions made from the beginning.

Europe's digital ambitions require more than technological innovation

Technological sovereignty is often associated with control over strategic technologies: processors, semiconductor manufacturing, communication infrastructure, cloud platforms, and AI models. For industrial Europe, however, sovereignty must also include the ability to integrate, operate, maintain, and evolve complex systems without becoming excessively dependent on opaque architectures, unsupported components or inaccessible engineering knowledge. Access to strategically important hardware does not by itself guarantee technological control over the complete system. An embedded platform also depends on development tools, in-house and third-party tightly coupled software components, or external services whose evolution and long-term support remain outside the system manufacturer's direct control. Architecture can limit some of these dependencies.

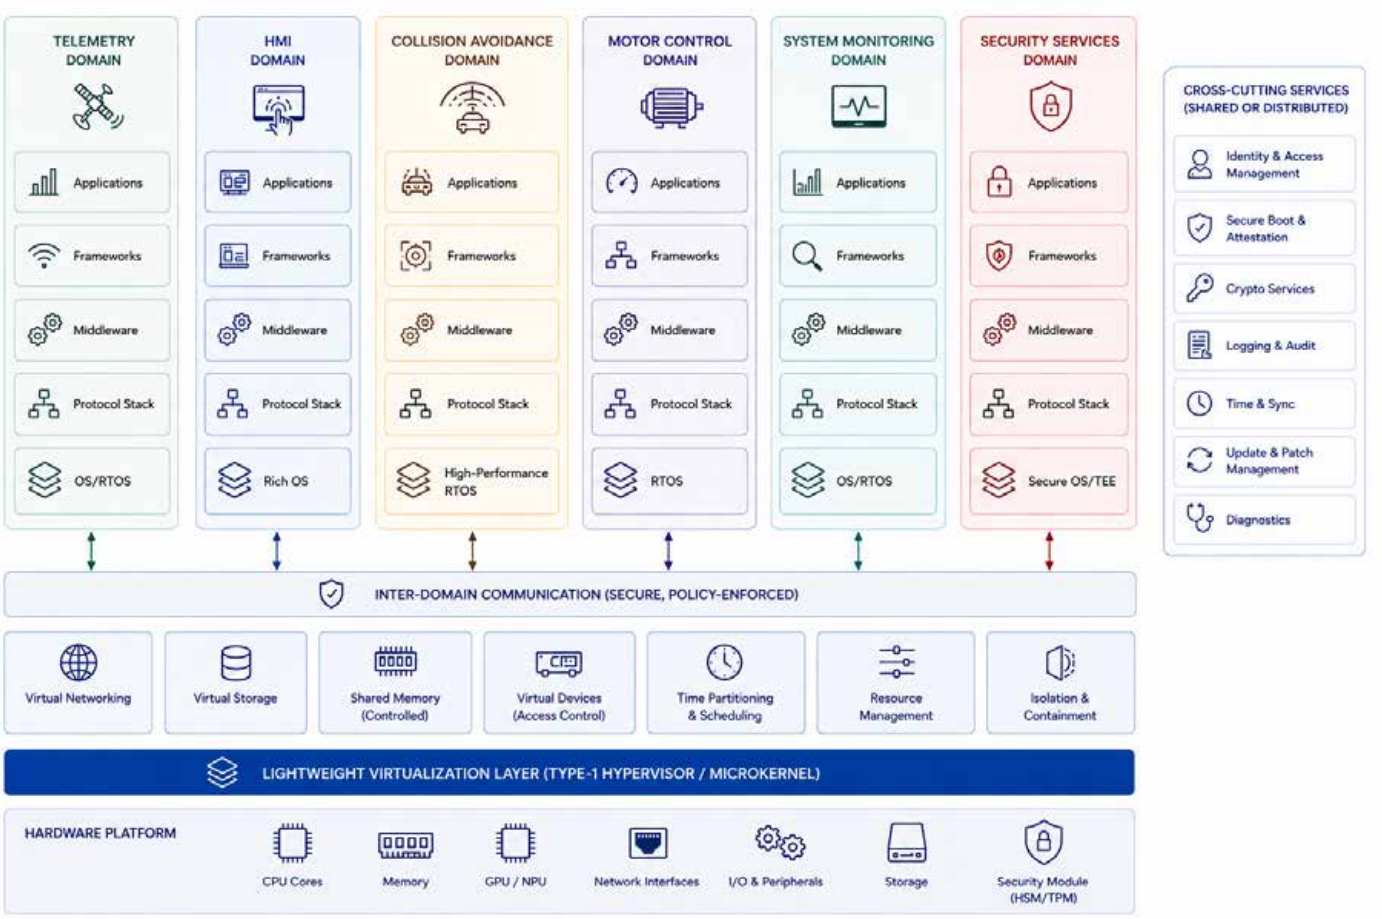

Stable interfaces, controlled software dependencies, and clearly separated functional domains can make individual technologies easier to replace without redesigning the entire system. This does not eliminate supply-chain risk, but it makes dependencies more visible and manageable. Modern embedded products typically combine software from several sources: operating systems, middleware, protocol stacks, open-source libraries, security components, diagnostic tools, and application software from specialised suppliers.
Each of these dependencies introduces assumptions concerning performance, timing, update availability, licensing, vulnerability management, and long-term support.

Technological sovereignty should therefore include the capacity to understand and control the entire systems created as a composition of components, system of systems. This requires investment not only in hardware and computing infrastructure, but also in software architecture, verification, lifecycle engineering, and the technical expertise required to maintain complex systems over time. These capabilities may be less visible than semiconductor fabs or AI supercomputers, but they determine whether technological investment can ultimately become dependable industrial infrastructure.

The growing complexity of modern embedded systems

The traditional embedded system was often designed around a relatively fixed set of functions. Software was closely coupled to hardware, connectivity was limited and substantial changes after deployment were uncommon. This model is rapidly changing. Embedded platforms are becoming connected, consolidated, and software-defined. Functions previously implemented on separate electronic control units are moving onto multicore processors and systems-on-chips. Cloud connectivity extends the system boundary beyond the device itself. AI workloads are increasingly executed at the edge, while software updates continue throughout the operational lifecycle.

These changes offer significant advantages, including hardware consolidation, improved diagnostics, remote maintenance, and faster functional evolution. However, they also create new forms of complexity. A safety-related control application may share processors, memory bandwidth, network interfaces or accelerators with less critical software. Diagnostic services may require privileged access to several subsystems. Update mechanisms become part of the security perimeter. AI applications may depend on external frameworks and data pipelines that evolve independently of the underlying device. This creates several engineering challenges for ensuring high levels of safety, security, and time-predictability.

Failures can propagate across domains that were previously physically separated. Consolidating workloads onto a common computing platform may reduce hardware complexity while making isolation significantly more important. Furthermore, predictable execution depends on more than the operating system. A real-time task may be scheduled correctly and still miss its deadline because another workload saturates memory bandwidth or a shared hardware resource.

In addition, software with different levels of criticality and assurance increasingly needs to coexist. A safety-related control function may operate alongside an open-source networking stack, a graphical interface, and an AI inference engine. At the design phase, treating all these components as though they

had identical security, safety, and lifecycle requirements is unsustainable. This is one reason mixed-criticality architectures are receiving increasing attention. They allow workloads to be separated according to their timing, safety, security, and assurance requirements while sharing a common computing platform.

The automotive industry provides an obvious example. Central computing architectures increasingly combine vehicle control, connectivity, infotainment, and advanced driver-assistance functions. Functional safety requirements defined through frameworks such as ISO 26262 coexist with cybersecurity engineering requirements such as ISO/SAE 21434. Comparable challenges exist in industrial automation, aerospace, railway systems, and medical devices. Complexity, in this context, is not simply the number of lines of code. It is the number of software components and their interactions, dependencies, and potential failure paths that engineers must understand and control. Working at the software architectural design level is becoming a de-facto solution to better control these aspects.

Software architecture as the foundation of trustworthy embedded systems

Trustworthiness does not mean assuming that a system will never fail. Faults, vulnerabilities, and unexpected conditions cannot be eliminated completely. Instead, trustworthy architecture aims to ensure that critical properties remain understandable and defensible when failures occur, software changes or attacks are attempted. Indeed, one of its fundamental principles is the definition of explicit trust boundaries. Every system contains components that must behave correctly for its critical properties to hold. These form its trusted computing base. The larger and more complex this trusted base becomes, the greater the number of components whose failure or compromise may affect the entire system. Reducing unnecessary privilege is therefore fundamental.

Memory protection, privilege separation, controlled interfaces, and least-privilege access can limit the impact of compromised or malfunctioning software. Hardware mechanisms such as virtualization, memory management units and controller I/O units provide useful building blocks, but they become effective only within a coherent architectural design. Spatial isolation prevents one workload from accessing memory, devices or data assigned to another. Temporal isolation aims to prevent one application from consuming resources in a way that compromises the timing requirements of another. Both are important in mixed-criticality environments.

Virtualization technology can provide mechanisms for partitioning workloads, if properly configured. The strength of the architecture also depends on shared resources, interrupt handling, memory behaviour, and the size and complexity of the privileged software layer. The same distinction applies to containers. Containers can provide useful software separation and deployment flexibility, but they generally share a common kernel and should not automatically be considered equivalent to stronger hardware-assisted partitioning.

In addition to the isolation, many software components with different criticality levels may communicate with each other to cooperate in accomplishing the same mission. As such, communication between domains must also be carefully controlled. Interfaces should be explicit, narrow, and independently testable. Data formats, access permissions, timing expectations, and failure behaviour should be treated as architectural properties. This limits the paths through which faults or attacks can propagate, also improving maintainability. Components connected through stable interfaces are generally easier to replace or update without affecting the entire system.

Once the software design is defined and developed, guaranteeing its integrity, and sometimes also its confidentiality, every time the device is powered on is crucial to ensure its properties and protect its assets. Secure boot, protected key storage, and device identity mechanisms can establish confidence in the software executing on a device. In distributed systems, attestation mechanisms may also be used to provide evidence about platform state. These mechanisms, however, are only parts of a broader security architecture. A correctly signed software image can still contain vulnerabilities, while weak key management can undermine otherwise strong cryptographic mechanisms. As such, software updates play a crucial role in the entire lifecycle of the device.

Connected embedded systems increasingly require security updates throughout their operational life. Update mechanisms must verify authenticity and integrity, handle interrupted installation and provide reliable recovery mechanisms. In critical environments, deployment may also require compatibility checks, staged roll-out, and operational approval. Being able to distribute an update is not the same as being able to activate it safely.

Furthermore, one important aspect is related to the software supply-chain visibility and controllability. Modern systems inevitably rely on open-source packages, third-party libraries, and externally developed firmware. Maintaining an accurate software bill of materials (SBOM) can help manufacturers identify components affected by newly disclosed vulnerabilities. Yet an SBOM is not itself a security assessment. Knowing that a vulnerable library exists in a product does not automatically determine whether that vulnerability is exploitable in the specific system architecture. Effective vulnerability management therefore depends on both software visibility and architectural understanding, that typically requires threat analysis and risk assessments at the design phase.

From regulatory compliance to engineering trustworthiness

European regulation is increasingly reinforcing a principle already familiar to experienced engineering organisations: cybersecurity cannot effectively be added at the end of product development. The Cyber Resilience Act introduces cybersecurity and vulnerability-handling obligations for products with digital elements throughout their lifecycle. Its reporting obligations start to become applicable in September 2026, while the broader set of requirements applies from December 2027. However, for manufacturers of long-lived embedded products, the architectural implications begin much earlier.

A system being designed today may still be operational in the 2040s. Decisions concerning update mechanisms, component traceability, cryptographic technologies, and software isolation therefore need to anticipate future maintenance requirements. As such, component selection itself becomes a lifecycle decision. Performance, purchase cost, and development effort remain important, but manufacturers must increasingly consider the availability of security updates, supplier vulnerability-management processes, and the possibility of replacing technologies whose support ends before the product itself reaches end of life.

In addition, the AI Act introduces other architectural challenges. As artificial intelligence becomes increasingly embedded within industrial and cyber-physical systems, engineers must manage the relationship between probabilistic AI components, and conventional deterministic software. An AI model used for perception, anomaly detection, or decision support may not behave in exactly the same way as traditional control software. This does not prevent its use in critical environments, but architecture must define its authority and its limits. An AI function may, for example, operate behind an independently designed monitoring mechanism, remain advisory rather than directly controlling an actuator, or be constrained by deterministic safety functions.

Functional safety, cybersecurity, and AI governance are separate disciplines with different objectives and evidence requirements, but they interact inside the same system. Architecture provides a common technical model through which these interactions can be understood and controlled. It identifies assets, trust relationships, communication paths, critical functions, and update boundaries. In doing so, it connects regulatory requirements to concrete engineering decisions. Compliance should therefore be seen as evidence that appropriate system properties have been applied and achieved, rather than as the primary purpose of engineering. A platform designed mainly to pass an assessment can still be fragile. A platform built around explicit boundaries, controlled dependencies, and predictable behaviour is more likely to remain trustworthy as both technology and regulation evolve.

Building embedded systems that remain trustworthy over decades

One of the most difficult challenges for embedded systems is the mismatch between industrial lifecycles and digital technology lifecycles. Industrial machinery, railway infrastructure, aircraft, or medical equipment can remain operational for decades. The technologies on which they depend, however, often evolve on much shorter commercial and support cycles. Semiconductor components may become difficult to source, while operating systems, development tools, and software libraries may no longer receive maintenance or security updates.

Architecture cannot eliminate this lifecycle mismatch, but it can reduce its impact by making critical dependencies easier to identify, isolate and, where necessary, replace. Hardware abstraction layers, stable interfaces, and portable application environments can reduce unnecessary dependencies on specific technologies. Abstraction should not become an objective in itself. Excessive abstraction can reduce performance and complicate timing analysis. The goal is to create flexibility where future change is likely while preserving predictability where hardware-specific behaviour is essential. Long-term maintainability also depends on preserving engineering knowledge.

Architectural decisions, security assumptions, timing requirements, and critical interfaces must remain understandable after the original development team has moved on. Configuration, build environments, and test evidence should remain reproducible wherever possible. Furthermore, a trustworthy architecture should also enable controlled evolution. Partitioning may allow a non-critical analytics application to be updated while a critical control function remains unchanged, provided that isolation mechanisms and interfaces remain valid. This does not automatically eliminate the need for impact analysis or recertification. It can, however, reduce the scope of change and make long-term evolution more manageable.

Cryptographic agility is another important consideration. Algorithms and security assumptions considered adequate today may become obsolete during the lifetime of a long-lived industrial platform. The ability to replace cryptographic mechanisms without redesigning the complete system is therefore an architectural advantage. Lifecycle engineering should also address system decommissioning. Embedded devices may retain cryptographic keys, credentials, configuration data, or operational information after they are removed from service. Their architecture should therefore support appropriate end-of-life procedures, including credential revocation, secure data removal, and the controlled disconnection of external services. Considering these requirements early in the design process helps to ensure that security and maintainability can be managed throughout the entire operational lifecycle of the system.

Software architectures as a European competitive differentiator

Europe's regulatory framework is sometimes described primarily as a constraint on innovation. Cybersecurity, lifecycle support, and functional safety requirements address

Key insights

- Software architecture contributes directly to technological sovereignty. Control over strategic technologies has limited value without the ability to integrate, maintain, and replace them.
- Mixed-criticality systems require more than software separation. Timing, memory, I/O, and shared hardware resources must also be considered.
- Cybersecurity is a lifecycle responsibility. Vulnerability management, updates, and software-component visibility need architectural support from the beginning.
- Compliance is an outcome of engineering. Regulations and standards define obligations, but trustworthy behaviour depends on architectural decisions.
- Maintainability contributes to safety and security. Systems that cannot be understood, updated, or migrated become increasingly difficult to defend over time.

genuine industrial risks. Connected devices operate in increasingly complex environments, software supply chains continue to expand, and critical products often remain operational long after their original technology has become obsolete. Responding to these challenges only through compliance activities performed late in development can increase costs without correcting fundamental design weaknesses.

An engineering-driven approach starts earlier. It establishes trust boundaries, controls communication, limits unnecessary privilege, separates workloads with different assurance requirements, anticipates updates, and technological change. These principles support not only compliance, but innovation itself. When critical functions are clearly separated, new software can be introduced with less disruption. When interfaces are stable, components can be replaced more easily. When dependencies are visible, supply-chain risks become more manageable.

Europe's technological competitiveness will therefore depend not only on the processors, AI technologies, and computing infrastructure it develops, but also on the quality of the systems built around them. Software architecture is the level at which technological innovation becomes dependable industrial capability, and where technological sovereignty becomes something that can be engineered, operated, and sustained.

Giorgiomaria Cicero
CEO & Co-Founder of Accelerat

Giorgiomaria Cicero is an Embedded Software Engineer, now CEO and Co-Founder of Accelerat, a deep-tech start-up company focused on advanced software solutions for next-generation cyber-physical systems like autonomous vehicles (e.g., automobiles, trains, drones, etc.), advanced robots, and factory automation. Giorgiomaria has an academic research background matured at the ReTiS Laboratory and experience in industrial R&D projects in real-time systems, virtualization technologies, and cyber-security solutions for modern heterogeneous embedded platforms.

About Accelerat

Accelerat is an Italian deep-tech company developing software technologies for safe, secure, and predictable cyber-physical systems. Its portfolio includes CLARE, a hypervisor-centric software stack, and the Bunkers family of secure edge-computing solutions through on-chip zero-trust architecture.

Among Bunkers, there are two pre-set variants: (i) Bunker for Linux that isolates critical Linux applications and OT operations from network-facing software, allowing protected functions to remain segregated even if the primary Linux environment is compromised; and (ii) AI Bunker that provides an isolated environment for protecting AI models at the edge against theft, reverse engineering, and unauthorized access while preserving standard inference workflows.

Trusted AI and digital sovereignty

The two fundamental requisites for scaling agentic AI

Chris Brockmann
CEO, Eccenca

Enterprises that describe themselves as AI-first are moving past pilots. They are no longer asking whether a chatbot can answer a customer question or whether a model can summarise a document. They are asking whether dozens or hundreds of AI agents can act across procurement, engineering, finance, and operations to ingest and interpret data, make recommendations, and act independently.

This is a different kind of undertaking than running a single application. It requires that every agent, regardless of which team built it or which large language model (LLM) powers it, operates on a shared and correct understanding of a business' context, knowledge and rules. And it requires that the organisation retains control over its data, its processes, and the meaning attached to both, even as the number of agents and use cases multiplies. These two requirements have names: trusted AI and digital sovereignty. Both are a matter of how an organisation chooses to manage its data and information in the long run. Get this wrong, and agentic AI fragments into inconsistent, unmaintainable point solutions. Get it right, and it becomes possible to orchestrate an entire ecosystem of agents with confidence.

What trusted AI requires

AI quality, i.e. hallucination vs (reproducible) accuracy, is often discussed in terms of training data volume. This misses the actual problem. AI applications and LLMs have enough data to dwell on. In fact, they have too much. The issue is that they understand neither the data's meaning nor the underlying specific business context. A large language model has no inherent knowledge of what "asset," or "revenue" means inside your organisation. It does not know that a "critical supplier" in procurement is defined differently than a "critical supplier" in risk management, or that a machine identified by one part number in an ERP system is the same machine referred to by a different identifier in a maintenance log.

LLMs work through probabilistic statistics: they predict likely relations and continuations of text and data points. Left to its own devices, an LLM will guess and will do this every time anew – plausibly, confidently and eloquently, though rather often wrongly. Multiply this across a landscape of agents built by different teams for different purposes, and you are heading for a new Babylonian era. Probability is not the same as correctness.

Trusted AI means:

- Every AI agent in the organisation shares the same understanding of business context regardless of which team deployed it or which data source it taps into.
- AI and LLM applications produce results that are reliable, reproducible and traceable, not just plausible and eloquently written.
- Models do not hallucinate facts about the business because they are not asked to infer them in the first place.
- Systems are guardrailed according to the organisation's actual rules and risk tolerance, not opaque security defaults.

To get dependable and reproducible outcomes, the probabilistic layer needs to be paired with something deterministic, a semantic layer that provides explicit meaning and context-sensitive rules for how data should be interpreted (Fig. 1).

Ontologies and knowledge graphs as the semantic layer

This semantic layer is what ontologies and knowledge graphs provide. In practical terms, an ontology defines the concepts that matter to a business – what a "supplier," an "asset," or a "work order" is, how these concepts relate to one another, and which rules govern them. A knowledge graph maps these semantics to the organisation's actual data, connecting information that lives in separate, often incompatible systems into a single, navigable network of meaning (Fig. 2). The effect is to take what is usually a Babylonian condition – dozens of systems, each with its own vocabulary, identifiers, and assumptions,

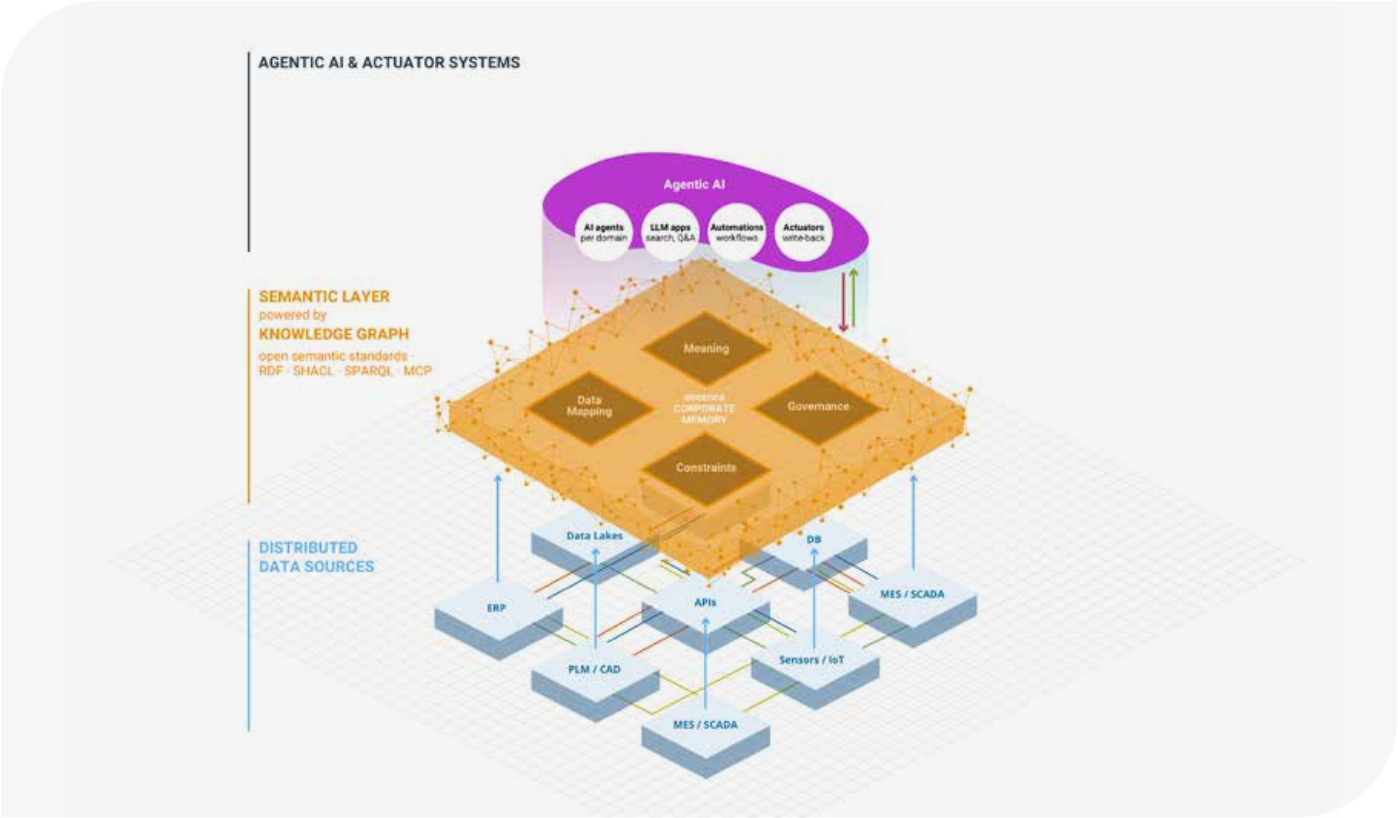

Fig. 1: The semantic layer sits between the agentic AI and the data sources.

built up over years of mergers, migrations, and departmental autonomy – and turn it into a consistent, unified understanding. The knowledge graph forms the semantic middle layer agentic AIs and LLMs must go through before even touching the data. Now every data point is no longer just a value in a table. It carries explicit business meaning and context that any agent, regardless of which model sits behind it, can read the same way.

Two Eccenca customers, who presented their journey to a harmonised cross-domain understanding and usability of their data on this year's Data Week Leipzig, illustrate what this looks like in practice:

- Cybersecurity at Airbus depends on unified cyber risk intelligence across a complex, multi-domain ecosystem with a highly diverse IT infrastructure. They use federated knowledge graphs and ontologies as an engine of certainty to create collectively consistent and trusted knowledge about their security posture.
- Austrian energy provider VERBUND faced a related but distinct challenge: harmonising data spread across non-compatible, historically grown systems in 113 hydro-power plants and 28 large dams, where the same information for an asset existed in different forms and cryptic abbreviations.

In both cases, the ontology and knowledge graph layer is what allowed them to work with dependable, shared meaning and data across different domains and infrastructure. This is the foundation trusted AI rests on. But building this layer raises an immediate follow-up question: who controls it, and what happens as the organisation wants to extend and scale it?

What does digital sovereignty entail in an AI ecosystem?

Like data and data sources, ontologies and knowledge graphs are tools that must be able to adapt to the dynamic changes an enterprise faces each day. Meanings, processes, rules and data sources might change. So can employed AI agents and use cases. In that context, digital sovereignty means your organisation retains genuine control over three things: your data, your semantic models, and the economics of changing both.

Control over data requires that everything can run on the organisation's own premises or its own cloud infrastructure, under its own governance and not exclusively inside a vendor's managed environment. It also requires that existing source systems do not need to be migrated or duplicated to be usable. Rather than moving data into a new proprietary store, interfaces are attached to the data where it already lives. This matters enormously in practice: large-scale data migrations are slow, expensive, and risky, and they create a new single point of dependency exactly where organisations are trying to reduce one.

Control over the semantic model requires transparency. The ontology needs to be understandable, not a proprietary black box inside a platform. Moreover, to retain the flexibility needed to adapt to dynamic changes, the ontology, the mapping logic, and the underlying data model should always belong to the enterprise, not a vendor's platform. This ensures that an organisation can expand its ontologies into new domains as the business evolves, combine multiple ontologies for cross-domain, cross-functional tasks, and reuse the same ontology across different projects and systems rather than rebuilding equivalent structures repeatedly. That's where open semantic standards such as RDF, OWL, SHACL and the broader W3C semantic web stack come into play, alongside emerging interoperability standards like MCP for connecting agents to data and tools. The use and support of open semantic standards are the foundation for flexibility and scalability needed in today's digital world.

Control over the cost of change is where sovereignty becomes a business factor. If every extension of an ontology or every reuse in cross-functional projects requires a multi-million Euro data management project, the cost of evolution multiplies with the number of changes an organisation needs to make. That number only grows as agentic AI spreads into more domains and use cases. Genuine sovereignty means the organisation's own teams have full control of their semantic layers, infrastructure and data-to-semantics-mappings. This is what separates a one-off implementation from a system that can keep pace with a business that keeps changing.

The combined requirement

Trusted AI and digital sovereignty should not be viewed as separate topics. Trusted AI depends on a shared semantic foundation across every agent in the organisation. Sovereignty determines whether that foundation can be built, controlled, and extended without becoming a new dependency and vendor lock-in. For leadership teams deciding how to scale agentic AI, the practical question is not which model or which vendor to select. It is whether the organisation is building a semantic foundation, i.e. an ontology and knowledge graph, that it genuinely owns: transparent, open standards-based, extensible and federate-able. That is what allows organisations to orchestrate an ecosystem of AI agents, not just a single application, with confidence and reliability across the entire business.

Chris Brockmann
CEO, Eccenca

Chris Brockmann is founder and CEO of the enterprise knowledge graph platform provider Eccenca GmbH and founder of the IT management consultancy brox IT-Solutions GmbH. His career has been shaped by a consistent focus on one challenge: enabling enterprises to turn increasingly complex and distributed data into reliable information, knowledge and ultimately better decisions – for humans and agentic AI applications.

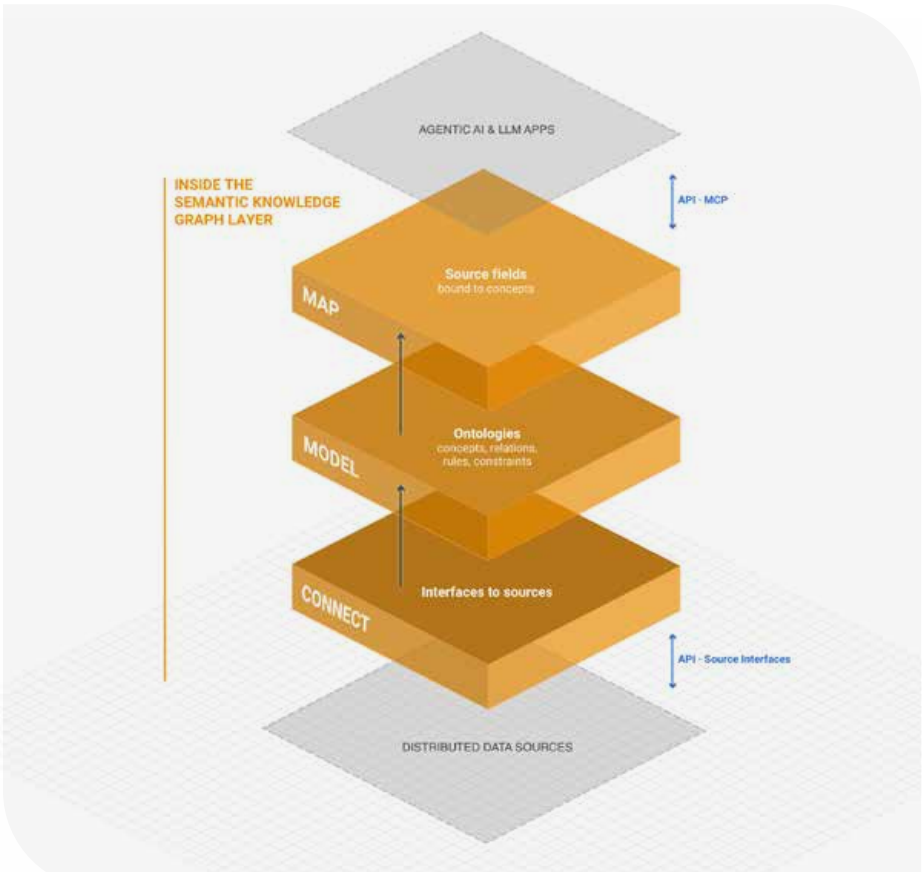

Fig. 2: The semantic knowledge graph layer provides meaning, rules and provenance to prevent LLMs from hallucinating.

New Member Focus

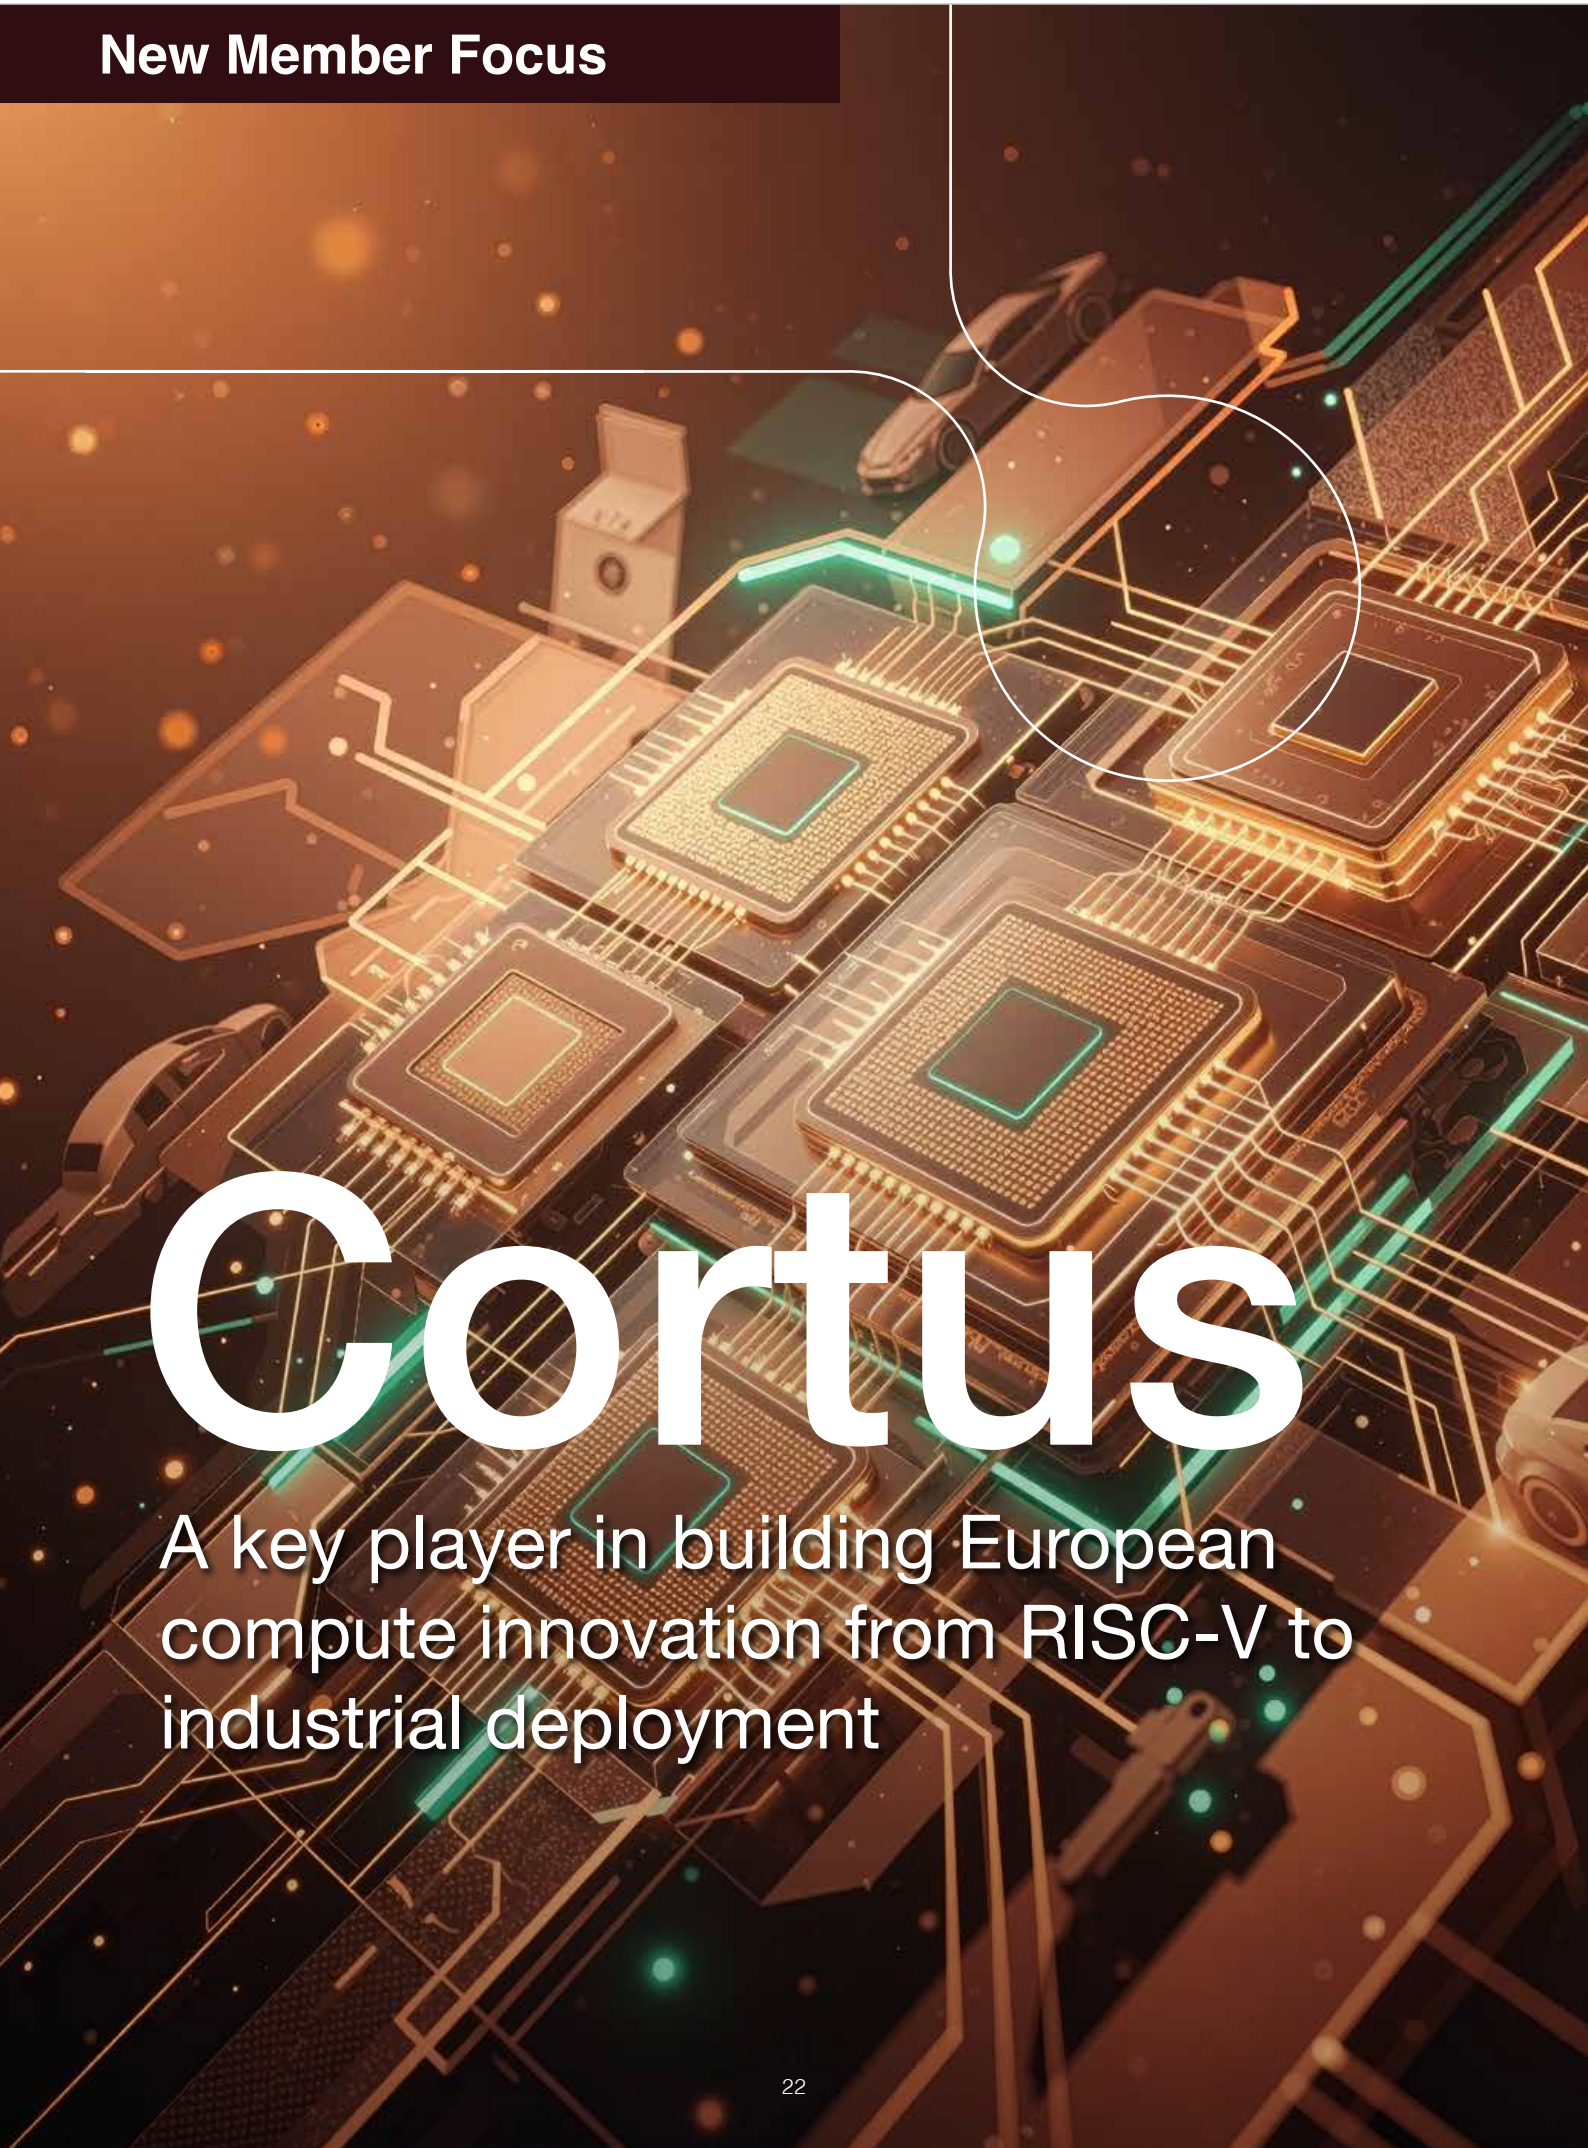

Cortus

A key player in building European compute innovation from RISC-V to industrial deployment

Michael Chapman
C.E.O. Cortus S.A.S.

Duc Nguyen
Cortus SAS

Cortus is a fabless semiconductor company specialized in the design and sale of RISC-V MCUs and AI inference chips for computer vision and image processing for the automotive, avionics, industrial and robotics markets. The company is one of the twelve founding members of the RISC-V Foundation and was the only non-American company among the original members at that time. It develops chips exclusively using its own processor IP, semiconductor technologies and design expertise, relying on a strong internal engineering team to deliver complete system-on-chip solutions.

A new era of intelligent, connected systems

As the digital transformation of industry is creating a new generation of intelligent systems, vehicles, factories, medical devices, energy infrastructures, and connected products increasingly require high-performance processing, artificial intelligence capabilities, security, and energy efficiency. This evolution is creating new challenges for the semiconductor industry. Innovation is no longer driven only by transistor scaling or manufacturing capacity. It increasingly depends on the ability to design efficient computing architectures, develop advanced semiconductor intellectual property and technologies, and transform research and engineering excellence into reliable industrial solutions.

As Europe strengthens its ambition for technological leadership and strategic autonomy, companies contributing to the semiconductor value chain play a critical role. This includes not only semiconductor manufacturing, but also processor architectures, IP development, system design, software ecosystems, and industrial deployment. As a new member of the INSIDE Industry Association, the company is committed to contributing its expertise in open architectures, embedded computing, and intelligent systems to a collaborative European technology landscape.

Two decades of open, embedded processor innovation

For more than two decades, its engineering teams have been developing processor technologies for embedded systems, with a strong focus on high performance, efficiency, flexibility, and scalability. The company has been a long-standing contributor to the open computing ecosystem and was among the early organizations involved in the development of RISC-V. The emergence of RISC-V represented an important shift in processor innovation, demonstrating that open architectures could become a foundation for industrial computing and enabling collaboration across companies, research organizations, and technology communities.

For the company, open architecture is not only a technical concept. It represents an approach to innovation based on accessibility, collaboration, and the ability to adapt computing solutions to specific application requirements. Over the years, its technologies have been integrated into embedded applications where reliability, low power consumption, security, and efficient processing are essential. These characteristics are particularly important in markets such as automotive, industrial automation, IoT, and intelligent edge devices.

Bringing artificial intelligence to the edge

Artificial intelligence is transforming the way electronic systems operate. While large-scale AI models have attracted significant attention through cloud computing, a growing number of applications require intelligence directly at the edge.

Edge AI brings computation closer to where data is generated. This enables:

- Faster response times
- Improved privacy
- Reduced communication requirements
- Lower energy consumption
- Greater autonomy for embedded systems

Deploying artificial intelligence at the edge requires a fundamentally different approach from traditional computing. Unlike cloud infrastructures, embedded systems must deliver high computational performance while operating within strict constraints on power consumption, silicon area, cost, reliability, and real-time responsiveness. Meeting these requirements calls for heterogeneous computing architectures, where general-purpose processors, AI accelerators, security engines, memory subsystems, and application-specific hardware work together seamlessly to deliver optimal performance and energy efficiency. Such architectures provide the flexibility needed to support a wide range of industrial applications while ensuring scalability, safety, and long-term reliability.

The company contributes to this evolution by combining expertise in processor architecture, AI inference acceleration, and semiconductor design capabilities. By leveraging its own processor IP and integrating specialized hardware accelerators into complete system-on-chip solutions, it enables energy-efficient intelligent computing platforms for demanding applications in automotive, avionics, industrial automation, robotics, and edge AI.

Enabling Europe's software-defined vehicles

The automotive sector represents one of the most significant transformations in computing history. Modern vehicles are evolving into software-defined platforms, combining advanced electronics, connectivity, artificial intelligence, and increasingly complex control systems.

Future automotive architectures will require computing platforms capable of supporting:

- Real-time control
- Artificial intelligence inference
- Cybersecurity
- Functional safety
- Efficient power management

While several automotive manufacturers are investing billions of euros to develop proprietary semiconductor solutions for their own vehicle platforms, Europe also needs technology providers capable of supporting a much broader industrial ecosystem. Building competitive semiconductor capabilities should not depend solely on vertically integrated companies, but also on the availability of European-designed reusable processor architectures, semiconductor IP, scalable computing platforms, and engineering expertise that can be shared across multiple industries.

This is where the company contributes to the European innovation landscape. By developing its own processor architectures, semiconductor intellectual property, AI acceleration technologies, and complete system-on-chip design capabilities, it provides scalable computing platforms that can be adapted to a wide range of applications, from automotive and avionics to industrial automation, robotics, and intelligent edge devices. Such reusable technology platforms enable European companies to accelerate innovation, reduce development time, and focus their investment on differentiated products rather than reinventing core computing technologies. Strengthening these foundational capabilities is an important step towards reinforcing Europe's competitiveness, technological resilience, and strategic autonomy in advanced semiconductors.

Collaboration across the innovation chain

The semiconductor industry is inherently collaborative. No single organization can address the complete technology stack alone. Progress depends on cooperation between semiconductor companies, research institutions, industrial users, software developers, and ecosystem organisations. This collaborative approach is one of the reasons the company joined the INSIDE Industry Association.

INSIDE brings together organizations working at the intersection of technology, research, and industrial innovation. The association provides a valuable platform for exchanging ideas, identifying new opportunities, and strengthening European cooperation in advanced technologies. By participating in the INSIDE ecosystem, it aims to contribute its experience in processor architectures, embedded computing, and semiconductor innovation while engaging with other European actors working on the future of intelligent systems.

Shaping the future of European computing

The future of computing will be shaped by the ability to combine openness, efficiency, security, and industrial relevance. Europe has significant strengths in research, engineering talent, and industrial expertise. Turning these strengths into global competitiveness requires strong collaboration across the entire innovation chain. From open processor architectures to AI-enabled edge systems, the next generation of intelligent devices will require new approaches to semiconductor design and deployment.

The company is proud to contribute to this European innovation journey and looks forward to collaborating with INSIDE members to transform advanced technologies into practical industrial impact.

Follow us on Social Media!

Let's get connected for our latest news, events and updates!

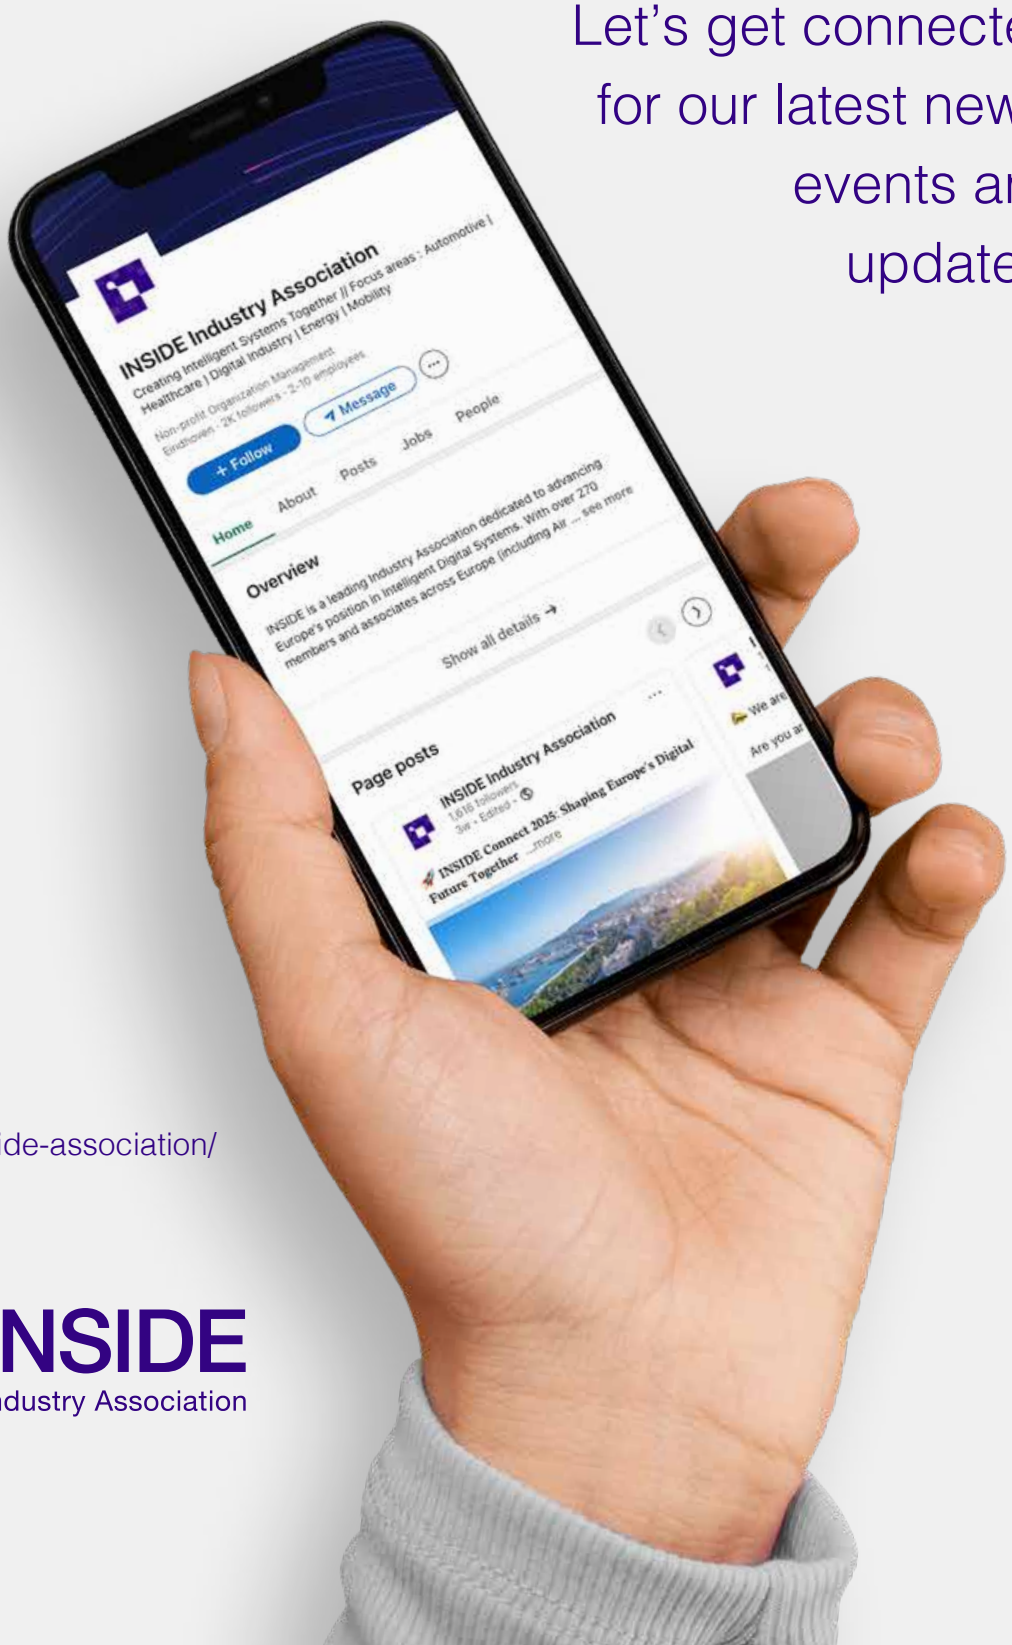

company/inside-association/

INSIDE
Industry Association

New Member Focus

Engineering trust from research to industrial deployment

How ResilTech turns research in safety, cybersecurity, AI and dependable computing into capabilities for Europe's critical digital systems

Francesco Brancati
Senior Director - R&D Innovation and Cybersecurity Business Development ResilTech S.r.l.

Andrea Bondavalli
Full Professor of Computer Science, University of Florence - Senior Scientific Advisor, ResilTech S.r.l.

Europe's digital transition is bringing semiconductors, embedded computing, artificial intelligence and connected infrastructure into a shared engineering landscape. Functions once implemented in isolated components are now distributed across software-defined platforms, edge devices, cloud services and operational environments. This convergence increases capability, but also dependencies, failure modes, attack paths and lifecycle complexity.

Europe's competitiveness will depend on its ability to deploy advanced technologies with confidence. This requires engineering practices that connect architecture, implementation, verification, validation, functional safety, cybersecurity and long-term operational support. Trust must be demonstrated through traceable requirements, evidence, controlled change and continuous risk management.

This capability can be described as engineering trust: the disciplined transformation of technical innovation into dependable industrial operation. ResilTech was founded around this challenge and has developed its expertise at the intersection of resilient computing, safety-critical engineering, cybersecurity and independent assurance.

From dependable computing to integrated resilience

Founded in Pontedera in 2007, ResilTech grew from dependable computing and safety-critical engineering. Its early work focused on automotive and railway systems, where hardware faults, software errors and rigorous assurance processes directly affect safety. As these systems became more connected and software-intensive, the company expanded into OT cybersecurity, critical infrastructure, embedded platforms and AI-enabled applications.

Today, more than 40 engineers and researchers work with customers and projects across Europe, Asia, North America and Japan. ResilTech specialists contribute to ISO working groups on automotive functional safety and cybersecurity, while its research teams participate in national and European programmes, including Chips JU initiatives. This multidisciplinary background enables ResilTech to operate across the engineering chain, from embedded technologies and product development to system integration and independent assurance.

Assuring AI-enabled critical systems

Machine-learning components introduce uncertainty that established verification processes struggle to capture. Their behaviour depends on training data, model

structure, operating conditions and the distribution of real-world inputs. Strong average performance can still coexist with unsafe misclassifications or unexpected responses in poorly represented scenarios.

ResilTech addresses this challenge through processes for the design, integration, verification and validation of AI/ML components. The work covers safety and security assurance, AI-specific threat modelling, benchmarking metrics and architectural mechanisms for detecting or controlling misclassification, aligned with ISO/PAS 8800 and emerging standards for safety-related artificial intelligence. Target applications include ADAS, automated driving, human-robot collaboration, cobots and industrial automation. In these domains, assurance requires evidence on relevant scenarios, known limitations, uncertainty, robustness and system behaviour when an AI component fails or returns an ambiguous result. European research projects and industrial activities allow these methods to be consolidated and transferred into reusable engineering and assurance capabilities.

Resilient embedded platforms for software-defined systems

Modern embedded platforms combine multicore processors, accelerators, complex SoCs and software-defined execution environments. Safety-critical applications must use these resources efficiently while retaining the ability to detect permanent hardware faults, timing degradation and aging effects during operation. Software Test Libraries are one enabling technology for resilient software-defined embedded platforms. An STL periodically exercises processor logic and checks its behaviour to detect hardware faults without dedicated duplicated hardware. ResilTech's current work includes STL development for RISC-V architectures, integration across bare-metal systems, operating systems and hypervisors, and hardware/software co-design to reduce execution overhead.

Fault coverage is a second research focus. The stuck-at model remains widely used, although advanced semiconductor technologies introduce effects that it captures only partially. ResilTech is studying path-delay faults and aging-related degradation, together with methods for assessing their impact on processor behaviour and diagnostic coverage. These activities connect semiconductor research with automotive, software-defined vehicles, industrial control and high-performance embedded computing.

Turning cybersecurity risk into engineering decisions

Cybersecurity requirements become useful when they are traceable to specific assets, functions, data flows and operational consequences. ResilTech develops systematic methods for establishing that connection across products, industrial systems and critical infrastructure. The process starts from architectural modelling. Data Flow Diagrams, TOGAF and ArchiMate models are used to represent assets, external entities, functions, data and trust boundaries. STRIDE supports threat identification, while exposure, attack feasibility and impact are evaluated to determine risk and derive attack paths toward critical targets.

For IEC 62443 applications, the resulting evidence supports Security Level Targets, zone and conduit design, defence-in-

depth measures, secure communication, access control, monitoring and verification. The methodology extends across the product lifecycle, from secure design and implementation guidance to vulnerability handling and update management. Applied across energy, railway, transport infrastructure and embedded products, this approach supports vector-based Security Level Targets under IEC 62443 and the derivation of technically justified requirements for programmes addressing NIS2 and the Cyber Resilience Act.

Research translated into industrial capability

ResilTech uses co-funded research projects as structured environments for developing methods, tools and specialist expertise for industrial programmes. The transfer follows a clear sequence: a recurring engineering problem is investigated, the resulting capability is consolidated, and the method is adapted to an operational context.

Approaches for deriving cybersecurity requirements for industrial control systems were developed through IESFactory and later applied in the energy sector. ArchiMate and TOGAF modelling matured through DistriMuSe and was subsequently used in transport-infrastructure cybersecurity activities. Research on threat modelling, secure development lifecycles, processor aging and path-delay fault models has also created opportunities for industrial uptake. The transfer includes people and organisational capability. Research projects train specialists, expose teams to emerging standards and generate reusable engineering assets that move into ResilTech's business units, supporting commercial work in cybersecurity, functional safety, embedded systems and assurance.

Leading Italy's contribution to Chips JU

Since 2024, ResilTech has held a continuous leadership role in Italy's participation in major Chips JU initiatives. The company acts as Italian lead in HAL4SDV and SHIFT2SDV and continues this responsibility in ODE4EC-DIG.

These initiatives connect semiconductor companies, embedded-software providers, automotive organisations, research centres and universities around next-generation digital platforms and European engineering capabilities. ResilTech contributes expertise in embedded systems, functional safety, cybersecurity, verification and assurance, while coordinating Italian partners within broader European programmes. This continuity strengthens relationships across the automotive and digital engineering value chains and creates a stable channel through which industrial needs inform research programmes and project results move toward market applications.

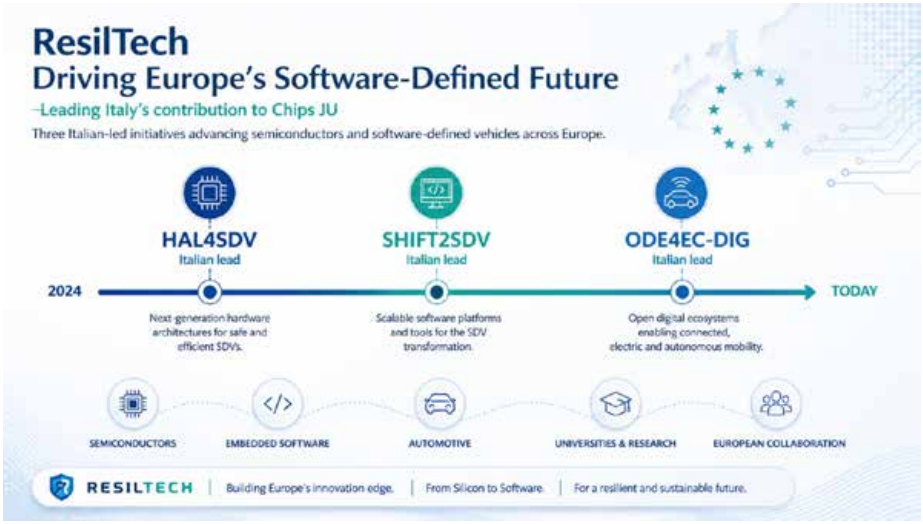

Engineering trust as a European capability

ResilTech works with universities and research centres through joint projects, training, thesis supervision and industrial PhD programmes. These relationships develop specialist skills and connect academic research with industrial engineering.

Looking ahead, engineering trust will increasingly have to address AI-native and agentic systems, where assurance, governance, constrained autonomy and human oversight become essential engineering capabilities, even though these areas are still emerging rather than mature industrial practice. This perspective aligns with INSIDE's mission to strengthen Europe's capacity to integrate semiconductors, software, artificial intelligence and cyber-physical systems into dependable and sustainable solutions. ResilTech contributes by converting collaborative research, standards work and industrial experience into capabilities that support Europe's critical industries.

Francesco Brancati
Senior Director - R&D Innovation and Cybersecurity Business Development ResilTech S.r.l.

Francesco Brancati is Senior Director at ResilTech, where he leads R&D innovation, co-funded programmes and cybersecurity business development. He holds an MSc and a PhD in Computer Science from the University of Florence. He coordinates ResilTech's role in European research initiatives, including Italy's participation in Chips JU projects, and has contributed to ISO working groups on automotive functional safety and cybersecurity.

Andrea Bondavalli
Full Professor of Computer Science, University of Florence - Senior Scientific Advisor, ResilTech S.r.l.

Andrea Bondavalli is Full Professor of Computer Science at the University of Florence and Senior Scientific Advisor at ResilTech. His research focuses on dependable and secure computing, safety-critical systems, resilient cyber-physical systems, trustworthy AI, and verification, validation and certification. He has coordinated major European research projects and has served in leadership roles in the international dependability community.

Bridging the gap between research and industrial deployment

Richard Hall-Wilton
Director of Sensors&Devices Centre of Fondazione Bruno Kessler - FBK

and Zahra Bisadi
Project Manager of Sensors&Devices Centre at Fondazione Bruno Kessler - FBK

Europe's ambition to strengthen its technological sovereignty in semiconductors depends on one critical challenge: transforming cutting-edge research into scalable industrial innovation. At Fondazione Bruno Kessler's Sensors and Devices Centre (FBK-SD), this transition is becoming a reality through a strategic combination of European initiatives, including IPCEI ME/CT and the Chips JU Pilot Lines. These flagship programmes are not only expanding FBK's research and technological capabilities but are also creating an innovation ecosystem where advanced semiconductor technologies can move efficiently from laboratory development to industrial deployment.

Building on decades of expertise in silicon sensors, micro- and nanofabrication, FBK is investing in next-generation platforms based on silicon carbide, silicon-germanium and 3D integration, supported by new cleanroom infrastructure and an open innovation model designed to foster collaboration with industry; this supports a vibrant lab-to-fab model. From high-performance sensors and MEMS for space, healthcare, environmental monitoring and autonomous systems to pilot manufacturing capabilities that accelerate technology transfer, these European projects demonstrate how coordinated investments in research infrastructures can bridge the gap between scientific excellence and market-ready solutions, reinforcing Europe's competitiveness in the semiconductor value chain.

Building the foundations of Europe's semiconductor future

The FBK Sensors and Devices Centre is built around a core mission of advancing sensor technology and achieving market adoption of the technologies. The centre has around 200 staff members including researchers, technologists, technicians, PhD students and interns. Its micro- and nano- facility, extending over 2000 sqm, is ISO 9001 and 27001 certified. The facility consists of 5 cleanrooms: CRD – silicon detectors, CRM

FBK-SD CR

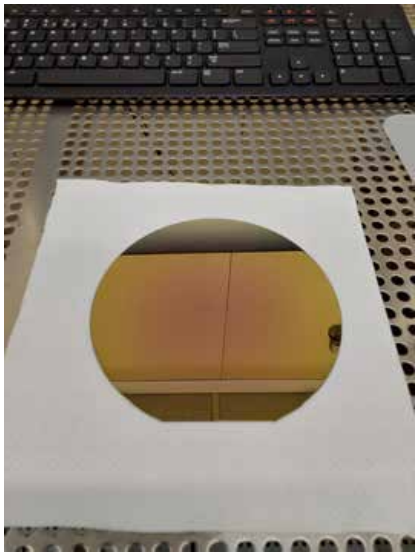

Amorphous Silicon Carbide layer deposited on silicon wafer by FBK-SD.

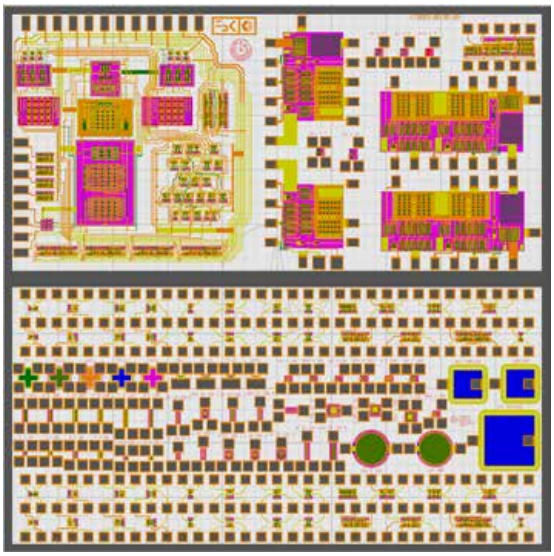

First tape-out of a SiC CMOS Readout ASIC by FBK-SD using Fraunhofer IISB technology.

– MEMS and NEMS, CRP – packaging and integration, CR3D – 3D integration, and CRSiC – SiC deposition (around 600 sqm currently under construction, thanks to IPCEI ME/CT). The cleanroom facility itself has been talked about in detail in issue 10 of the July 2025 issue of INSIDE Magazine.

With a decadal investment budget of around 100 million euro, FBK-SD is participant in IPCEI ME (closed in June 2025), IPCEI ME/CT, and 4 Chips JU Pilot Lines: Wide Bandgap (WBG), PIXEurope, SUPREME and DIREQT.

IPCEI ME/CT is a collaborative initiative aimed at driving innovation, research, and industrial deployment across the microelectronics and communication technologies value chain. Led by fourteen Member States of the European Union, this project represents a strategic endeavor to foster digital and green transformation while enhancing Europe's competitiveness in key sectors

Utilising the technologies built on the success of the IPCEI ME project, the Centre is actively driving its proprietary R&D roadmap within IPCEI ME/CT Sense Workstream, focused on advancing microelectronics technologies to produce advanced sensing systems that feature high precision, low energy consumption, embedded intelligence, sensor fusion capabilities, and high reliability applicable to various sectors, ultimately supporting Europe's technological and industrial leadership in semiconductors development.

In IPCEI ME, FBK-SD developed the 3D integration platform applied to SiPM (Silicon photon multipliers) and SPAD (Single-Photon Avalanche Diode) technologies and in IPCEI ME/CT, it is extending its micro- and nano-fabrication capabilities – both infrastructure and equipment – to Group IV semiconductors other than silicon, particularly SiC, as well as to develop new hetero-integration capabilities to provide advanced silicon-on-chip/silicon-on-package solutions. Silicon carbide has well-known advantages for power electronics; these advantageous properties can also be key in advancing sensor technology and capability, though silicon carbide sensors are at a very nascent stage of research currently. During the first two years of the project, significant progress has been achieved in establishing SiC technologies. Our team successfully deposited amorphous SiC on silicon substrates and completed the first tape-out of a SiC CMOS readout ASIC using Fraunhofer IISB technology.

Experimental results from the amorphous SiC deposition and surface characterisation closely matched theoretical predictions, confirming the validity of the developed approach. The project has also strengthened collaboration with European suppliers,

SiC diode post-lithography inspection image.

supporting both the procurement of advanced equipment, the co-development of a SiC processing system and material evaluation as part of process research.

Accelerating the path from research to industrial impact

FBK is playing a key role in Europe's effort to strengthen its semiconductor ecosystem through its participation in four Pilot Lines co-funded by the Chips JU. Access to pilot, small and medium fabrication capacity is the main barrier to market entry of novel devices. These large-scale European initiatives are designed to bridge the gap between research and industrial manufacturing, accelerating the uptake of strategic semiconductor technologies while reinforcing Europe's technological sovereignty, innovation capacity, and supply chain resilience. By contributing its expertise in micro- and nano-fabrication, photonics, quantum technologies, and advanced materials, FBK is helping transform cutting-edge research into manufacturing-ready solutions for the next generation of electronic and quantum devices.

As part of the WBG PL consortium, FBK is contributing to the development of wide bandgap (WBG) semiconductor technologies specifically for non-power applications. FBK is realising advanced sensor technologies particularly for harsh chemical, physical, high temperature, or radiation environments, and RF-MEMS electronics. It is also looking at SiC waveguides as part of its photonic integrated circuits in the NIR-VIS regime, and defect engineering in silicon carbide for quantum and photonic applications. By strengthening the capabilities of the participating pilot-line facilities, the initiative addresses key front-end manufacturing challenges and enables the development of high-performance devices that surpass the limits of conventional silicon technologies.

Within PIXEurope, FBK-SD joins 19 other institutions from 11 European countries to establish a globally unique Open Access Pilot Line for Photonic Integrated Circuits (PICs). The project addresses critical technological and manufacturing challenges while supporting RTOs and industry in accelerating the transition from laboratory innovation to industrial production.

As a partner in the quantum pilot line SUPREME, the SD Centre is helping bridge the gap between cutting-edge research and industrial-scale manufacturing of superconducting quantum technologies. Through pre-commercial fabrication services for superconducting quantum devices, the pilot line enables researchers and companies to design, prototype, and validate innovative solutions using advanced, standardised manufacturing processes, accelerating their path towards market-ready applications.

FBK is also contributing to DIREQT, the European pilot line dedicated to diamond quantum technologies. Covering the entire diamond value chain, the initiative aims to transform laboratory advances into scalable manufacturing, paving the way for the large-scale production of diamond quantum chips and supporting the emergence of a competitive European quantum ecosystem.

In summary

The adage of the enthusiastic investigative journalist is *"follow the money"*; this is very apt for the deeptech semiconductor ecosystem too, where huge investments are needed to realise technological progress. As shown above, the current investments in IPCEI ME/CT and in the Chips JU Pilot Lines are enabling future products to bridge the gap between research and industrial deployment, mediated by research and technology organisations with a specific remit in applicative development. This investment level needs to be sustained and intensified if the ambition for Europe to produce a higher fraction of its own chips is to be realised. Similarly, there is a contrast between the speed of research and that of the semiconductor market; again, this is best bridged by RTOs with a focus on getting novel innovation out there in the market.

New Member Focus

Malta Enterprise

MALTAENTERPRISE

Building an innovative ecosystem for Europe's next chapter

For a small European economy, competitiveness is not simply about attracting investment. It is about creating an environment where businesses can establish, grow and innovate to international markets. With more than five decades of supporting Malta's economic transformation, Malta Enterprise has evolved alongside the country's economy and the changing dynamics of global competitiveness. Its role today reflects this evolution: from supporting enterprise and investment to encompass innovation, technology, research and development, skills, internationalisation and the development of new economic capabilities.

Malta Enterprise places sustainable development and quality of life at the heart of its approach to growth, recognising that long-term competitiveness depends on strengthening economic resilience, developing new capabilities and creating opportunities for people and businesses. The ambition is clear: to strengthen Malta's existing economic base while creating the conditions for new technologies, sectors, and business models to emerge.

From investment support schemes to ecosystem development

Malta Enterprise administers a broad range of support measures covering investment, business development, innovation, research and development, and skills. These instruments can help companies expand their operations, adopt new technologies, strengthen their capabilities and move towards higher-value activities.

However, successful investment depends on more than just incentives. A strong business environment is built on a range of interconnected factors including talent, infrastructure, research, technology, finance, networks and access to markets. This is why Malta Enterprise is increasingly focused not only on individual support measures, but also on strengthening the connections between these different elements and creating an environment in which businesses can grow, innovate and compete internationally.

Striking the right balance between strengthening existing capabilities and

creating the conditions for new ecosystems to emerge, is particularly important as Malta builds on its strong economic performance. With Malta among Europe's fastest-growing economies, the next challenge is to ensure that growth is increasingly driven by innovation, productivity, skills and higher-value activities, thereby enhancing the economy's resilience and competitiveness over the longer term.

Policy shaped by industry

Achieving this requires more than financial incentives. Economic policy must also reflect the conditions businesses actually face in practice and respond to the opportunities and challenges emerging across different sectors. Malta Enterprise plays an important role at the interface between government, industry and the wider stakeholder community. Engagement with businesses, industry representatives, educational institutions and other stakeholders provides valuable insight into investment barriers, infrastructure needs, skills gaps, technological developments, regulatory considerations and emerging market opportunities.

This dialogue helps ensure that these realities are taken into account when policies and support measures are designed and refined. It also creates an ongoing feedback loop, allowing the experience and the needs of enterprises to help shape future policy.

In this sense, Malta Enterprise's role extends beyond the implementation of support measures. It also contributes to shaping an economic environment in which businesses can respond to changing opportunities, build capabilities and remain competitive over the longer term.

The Malta Semiconductor Competence Centre: an ecosystem in practice

The Malta Semiconductor Competence Centre (MSCC) provides a tangible example of this approach.

Led by Malta Enterprise, the Centre brings together academia, research and industry around semiconductors, a field of growing strategic importance to Europe's technological resilience. Its activities encompass skills, research, entrepreneurship and international collaboration, connecting different parts of the innovation ecosystem.

The MSCC also demonstrates the importance of strengthening existing capabilities while developing new niches and international connections. Through programmes such as ChipStartEU and partnerships with international semiconductor organisations, Maltese and European start-ups and companies can gain access to expertise, networks and opportunities extending well beyond the domestic market.

The objective is therefore not simply to establish a new initiative, but to create the capabilities, relationships and talent base from which a sustainable sector can develop.

These efforts sit within the wider direction of Malta Vision 2050, which provides a long-term framework for sustainable economic development, resilience, quality of life and competitiveness.

For Malta Enterprise, the challenge is to connect this long-term ambition with decisions being made today: how existing companies innovate, where new investment is attracted, which skills are developed, which technologies are adopted and how Maltese enterprises connect with global markets.

The broader European lesson is clear. Competitive innovation ecosystems are not created by a single programme or institution. They emerge when policy, industry, research, skills, investment and international networks reinforce one another.

Malta Enterprise's evolving role reflects this interconnected approach: supporting existing businesses to grow and innovate, creating space for new niches and ecosystems, listening to industry and stakeholders, and connecting Malta's capabilities with international opportunity.

In doing so, it can help turn sustainable development and quality of life into genuine drivers of growth, while strengthening Malta's contribution to a more innovative, resilient and globally connected Europe.

About Malta Enterprise

Malta Enterprise is Malta's economic development agency, supporting the growth of existing operations while attracting new foreign direct investment. It also acts as an adviser to government on economic policy, drawing on its close interaction with Malta's economic stakeholders.

Connect with Malta Enterprise
www.maltaenterprise.com
LinkedIn: https://www.linkedin.com/company/malta-enterprise

Discover Malta Enterprise's latest initiatives, support measures, and opportunities for international collaboration.

Chips JU

HECS 2026

2&3 December
Helsinki

SAVE THE DATE!

Aeneas

EPoSS. European Association on Smart Systems Integration

INSIDE Industry Association

Research Project Highlight

Dare to share

The value of chips from smart cities to collective urban intelligence

Alper Kanak, Ph.D.
NexTArc Technical Coordinator / R&D Director

Baran Çürüklü
NexTArc Project Coordinator / Senior Lecturer

Charlie Gullström
Dissemination, Communication and Exploitation Workpackage Leader [BC1.1]/ Head of Research and Innovation, Adjunct Professor in Circular Economy for Architecture and Planning

Matthias Haase
Scientific and technical dissemination responsible [BC1.1]/ Professor

What comes after the Smart City? This question lies at the heart of NexTArc — Next Generation Open Innovations in Trustworthy Embedded AI Architectures for Smart Cities, Mobility and Logistics. Bringing together 38 partners from across 10 European countries, NexTArc investigates how embedded Artificial Intelligence (AI), trusted electronics, edge-cloud computing, mobility, logistics and urban infrastructure and design can converge to make European cities and industrial ecosystems more liveable, sustainable, resilient and climate-neutral. Besides trying to create synergies through these technologies the project underlines the importance of open solutions as well as being pragmatic, thus assume the stakeholder perspective.

For more than two decades, Europe has invested in smart buildings, connected mobility, intelligent energy systems, IoT platforms, digital twins, edge computing and AI. Yet many cities remain collections of independently smart systems rather than genuinely intelligent ecosystems. One major problem is the lack of holistic view and understanding of the importance of long-term commitment by the decision makers. The industry on the other hand may not be decisive enough regarding understanding the need of different stakeholders. The latter requires the focus switch from a bottom-up technology push to a top-down stakeholder focused technology development. Thus, the next step both for the NexTArc project and the European industry is therefore not simply to deploy more sensors, or to create another data platform, or to introduce another AI model. The challenge is to assume the real needs of the citizens, align the industrial ecosystem, finally connect existing and emerging capabilities over i.e., buildings, vehicles, energy systems, urban infrastructure, and their users.

This was a key take-away from the NexTArc Innovation Workshop on 8–10 June 2026, hosted by SWECO as part of the EU New European Bauhaus Festival in Brussels, in which NexTArc partners and participants shared their experiences from complementary European projects and initiatives in three interconnected domains: Circular Cities and Regions, Energy-Efficient Cities and Communities, and Sustainable Mobility and Transport. One memorable workshop expression was “future smart buildings of the past.” IoT, Big Data, AI, smart buildings, blockchain and edge computing are no longer innovations simply because they exist individually. The opportunity lies in combining them into trustworthy applications and services that could previously not be deployed efficiently or at scale.

This leads to a different question: what becomes possible when Europe’s digital, physical and industrial assets can safely share not only data, but also knowledge, capabilities and intelligence, based on a willingness to share (at local city district level)? This is the proposition behind Dare to Share.

NexTArc: turning technologies into systems

NexTArc draws on the premise that European cities cannot solve their challenges within the respective boundaries that individual technologies, sectors, municipalities, regions or countries have formalised. Instead they need to be interconnected to facilitate a cross-fertilization to foster the innovation potential. The need for cross-fertilization can be illustrated in many ways: Mobility affects energy demand. Transport connects urban and industrial spaces. Logistics influences resource efficiency and competitiveness. Embedded AI enables new services but simultaneously introduces requirements for cybersecurity, trust, computational efficiency and dependability. A value-based and citizen-

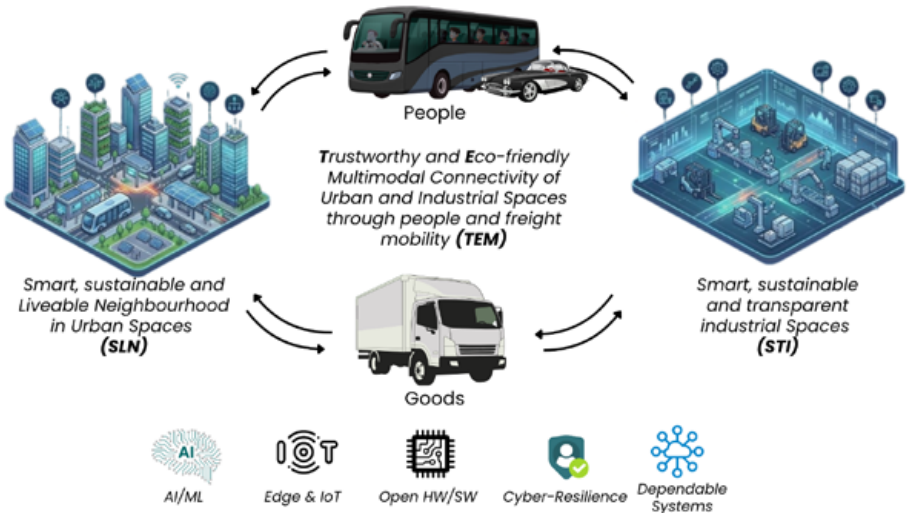

centred / actor-focused approach is therefore necessary in order to develop trustworthy multimodal and integrated technologies to an urban and industrial system.

Technically, the project addresses this convergence through a five-fold innovation approach: cyber-resilience on chip, low-power embedded AI, heterogeneous edge-cloud connectivity, dependable high-performance computation, and trustworthy services. These capabilities span the computing continuum, from individual components and embedded devices to vehicles, edge systems and complex mobility, logistics and urban applications.

NexTArc's innovation model converges in three complementary use cases. Smart, Sustainable and Liveable Neighbourhoods in Urban Spaces (SLN) integrates services including indoor food production, energy and water reuse, multi-vehicle charging, and air and noise monitoring. Smart, Sustainable and Transparent Industrial Spaces (STI) addresses trustworthy cross-organisational supply chains and internal logistics. Trustworthy and Eco-friendly Multimodal Connectivity of Urban and Industrial Spaces (TEM) connects these environments through more efficient movement of people and goods.

These are not three isolated demonstrators. They represent interacting urban spaces, industrial spaces, people and goods, supported by the five-fold innovation approach. NexTArc therefore raises a question of wider European relevance: not simply how can individual technologies become smarter, but how can they become trustworthy systems that create societal and industrial value together?

Dare to Share: from data sharing to intelligence sharing ("willingness to share")

"Sharing" should not mean placing every piece of urban or industrial data into a central cloud. Rather, the value of sharing should be addressed with a system's approach at the local level where European actors operate. Europe's strengths in privacy, cybersecurity, embedded electronics, edge computing and industrial engineering enable a more distributed model: data can remain close to where it is generated while authorised actors access the knowledge, predictions or services derived from it.

Early-on and in line with the above, the NexTArc system architecture discussions consequently moved away from the idea of a monolithic platform or "digital twin that does everything." Instead, the emerging approach connects global data, sensing, edge processing, fog computing, integration/data hubs and targeted applications. Semantic descriptions, access control, authentication and metadata can enable interoperability while computational tasks are executed where appropriate.

As an example, let's put forward a municipality that needs congestion predictions. At the same time a local energy operator needs forecasts of charging demand. A facility manager simultaneously requires an agent to monitor abnormal consumption. And a logistics operator needs to determine electric-fleet missions can be completed on time. These actors do not need the same application. They need trusted access to interoperable information and capabilities relating to the same geographic location. Whether an urban or industrial system, their willingness to share (data and other resources) will depend on their trust in others in the local community.

This points towards Community Hubs for Agentic Insights and Sharing: local or federated environments where citizens, municipalities, infrastructure operators, SMEs, researchers and industries can turn distributed information into useful decisions. NexTArc use cases are all small-scale community settings with clear boundaries where sharing of resources and data enables and accelerates system transformation to circular economy, improved logistics and production.

Three domains, one ecosystem
Circular Cities: from waste management to resource intelligence

A circular city requires more than recycling. It needs to identify which resources exist, where they are, how they are being used and by whom; and how they can re-enter productive value chains to avoid any waste from the system. Digital Product Passports, digital twins, semantic models and lifecycle information can transform physical assets into discoverable resources. Buildings (building materials and components), batteries, vehicles, electronics and infrastructure could carry machine-readable histories describing composition, condition, maintenance and reuse potential.

This leads towards circularity-by-design.

Interoperability does not necessarily require everyone to use exactly the same model. Different organisations can retain their own information structures while semantic mappings establish relationships between them. The longer-term opportunity is to transform cities from linear consumers of resources into resource-intelligence ecosystems.

Energy-Efficient Communities: intelligence where energy is consumed

AI can optimise energy, but AI itself consumes energy.

Informed by the Brussels discussions the Dare to Share model therefore identifies energy-efficient edge AI as a key direction: intelligence operating locally on low-footprint computing hardware so that the technologies supporting sustainable cities do not themselves become an unsustainable burden. The emerging progression starts from the devices and extends to buildings, districts and finally the community.

Buildings evolve from passive consumers into active stakeholders with needs that encompass photovoltaics, storage, HVAC, heat pumps and EV charging with predictive control. The future resource- and energy-efficient city therefore does not simply generate more renewable electricity. It becomes better at coordinating when, where and why energy is consumed.

Sustainable mobility: the digital vehicle of the future

Connected, Cooperative and Automated Mobility (CCAM), software-defined vehicles, intelligent batteries, charging infrastructure, Vehicle-to-Grid and digital twins increasingly connect transport with the city's digital and energy infrastructure.

The Digital Vehicle of the Future should therefore be understood not merely as a transportation asset, but as a mobile sensing, computing, energy and intelligence node.

This also opens the door to citizen-centric CCAM co-design. Citizens should not remain passive consumers of mobility systems. Participative and collaborative monitoring could provide a richer understanding of the "mobility status" of the city, combining transport performance with accessibility, safety, health, environmental impact and affordability.

Future mobility quality is becoming a function of:

Mobility Quality = f(Time, Health, Cost, Safety, Energy, Eco-friendliness).

The goal is not merely faster mobility, but better mobility for people and cities.

Cross-fertilising Europe's innovation assets

NexTArc's cross-fertilisation ecosystem connects lessons from other EU projects, e.g., OPEVA on electric-vehicle autonomy and battery intelligence; AI4SWEng on AI-supported software engineering and Small Language Models; ESCALATE on zero-emission heavy-duty mobility and logistics; TEAMING on physics-informed predictive powertrain health; and VALU3S on verification and validation. Complementary initiatives such as AIonBMS extend this discussion towards intelligent energy-management. Participants' wider innovation trajectory—including SCOTT, AI-TWILIGHT, PowerizeD, SUSTRONICS, AIMS5.0, ALL2GaN and LoLiPoP IoT—brings experience ranging from trusted IoT to intelligent lighting. The urban and energy perspective is further strengthened by experiences from PEDvolution, addressing interoperable Positive Energy Districts; UPCHANGE, focusing on neighbourhood-scale clean-energy transition; and CABE, on secondary-resource sharing.

Together with strategic directions emerging through INSIDE, Chips JU, ECAVA, 2ZERO, CCAM, the Cities Mission, CETP and DUT, the challenge is to cross-fertilise these initiatives across programme and sector boundaries, integrate them into open and trustworthy systems, validate them in real environments, and scale them into shared European solutions for circularity, energy, mobility and industrial competitiveness.

Three principles for Europe's urban future

Across these discussions, three design principles emerge.

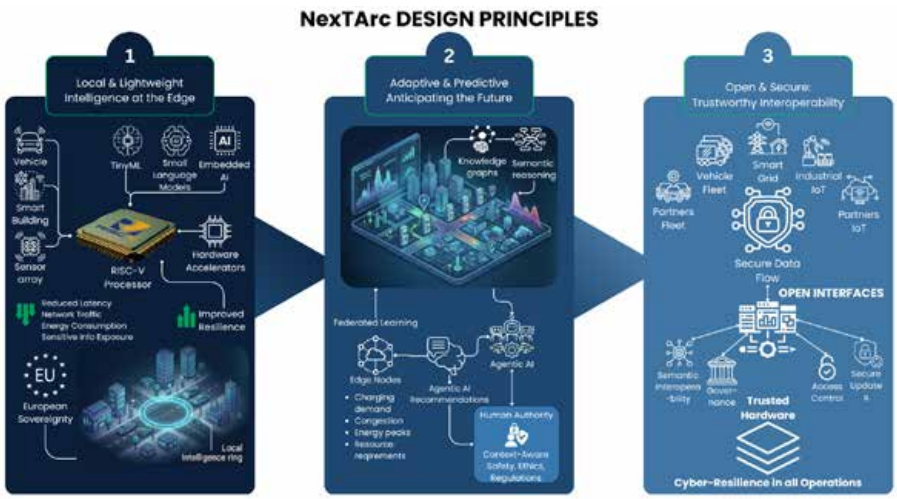

Local & Lightweight

Not every decision belongs to a hyperscale cloud. Local processing can reduce latency, network traffic, energy consumption and exposure of sensitive information while improving resilience. European sovereignty will depend not only on AI models, but also on the processors, accelerators and trusted electronics on which intelligence executes.

Adaptive & Predictive

Future urban infrastructure should not merely report what has happened. It should anticipate what is likely to happen next. Digital twins, knowledge graphs, semantic reasoning, federated learning and agentic AI can transform observations into context-aware recommendations and predictions of charging demand, congestion, energy peaks, infrastructure degradation or resource requirements, to monitor, to understand, to predict, to coordinate, and to adapt.

Open & Secure

Open interfaces cannot mean uncontrolled interfaces. Future ecosystems require semantic interoperability, governance, trusted hardware and cyber-resilience. They also require certifiability-by-design: systems engineered to continuously generate trustworthy evidence that can support assurance, conformity assessment and, where applicable, future formal certification, without confusing R&D validation with official certification.

Community Hubs: making intelligence useful

Community Hubs should not become another name for urban data platforms. Their purpose is to provide trusted meeting points between people, infrastructure, data, models and AI agents.

For citizens, they can translate complex information into understandable insights. For municipalities, they can provide predictive views of mobility, energy and infrastructure. For SMEs, common interfaces can lower barriers to creating new services. For industry, they can generate markets around embedded AI, semiconductors, digital twins and trusted data services. The key lesson is simple: users want useful applications and outcomes, not platforms.

From "more chips" to "innovation with chips"

Society does not consume chips as an end in themselves. It consumes mobility, energy,

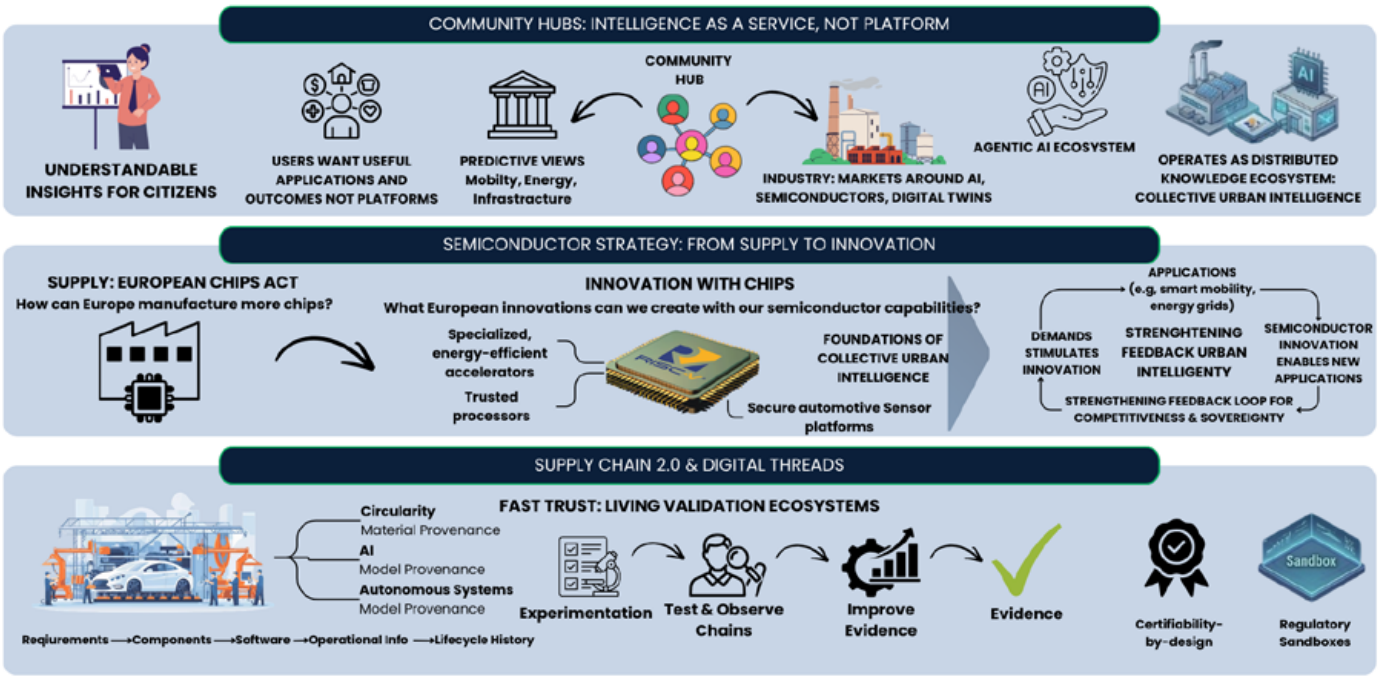

healthcare, industrial productivity and digital services enabled by chips. Europe should therefore complement:

"How can Europe manufacture more chips?"

with:

"What European innovations can we create with our semiconductor capabilities?"

The cities begin to operate as distributed knowledge ecosystems, Collective Urban Intelligence. European innovations can become foundations of Collective Urban Intelligence. Applications create demand; demand stimulates semiconductor innovation; semiconductor innovation enables new applications. Strengthening this feedback loop is essential for European competitiveness and sovereignty.

Supply Chain 2.0 and Fast Trust

Future innovation ecosystems must manage not only physical products, but continuous flows of data, models, software, knowledge and evidence. A software-defined vehicle is never truly finished when it leaves the factory. Software evolves. AI models evolve. Cybersecurity threats evolve. Batteries age. Regulations and infrastructure change.

Supply Chain 2.0 therefore connects requirements, components, operational information, lifecycle histories and evidence through digital threads. Europe needs faster mechanisms for moving innovations into realistic environments. The objective is not weaker regulation, but a shorter distance between technological iteration and trustworthy deployment. This calls for Living Validation Ecosystems, certifiability-by-design, digital evidence chains and regulatory sandboxes where emerging systems can be tested, observed, improved and progressively evidenced.

The result is Fast Trust: rapid experimentation accompanied by continuous evidence generation.

Dare to Share: from NexTArc to a European agenda

NexTArc illustrates why cross-fertilisation matters. Its use cases connect urban spaces, industrial environments, people and goods; its innovations connect electronics with AI, edge-cloud systems, cybersecurity and dependable computing; and its workshops connect these developments with a wider European R&I ecosystem.

Not one central intelligence controlling the city, but many local intelligences cooperating securely and human-centrically. This is the essence of Collective Urban Intelligence—and it gives Europe's semiconductor, AI, mobility, energy and urban strategies a shared direction: Local & Lightweight. Adaptive & Predictive. Open & Secure.

The next generation of European innovation may depend less on asking what else we can connect and more on deciding what we can safely share, what intelligence we can collectively create, and how that intelligence can improve everyday life, competitiveness and sovereignty. NexTArc provides an important European environment for exploring that transition: from components to systems, from research to application, and from isolated innovation to shared progress.

What's next: from foresight to a European innovation track

The ambition of NexTArc is to ensure that the Dare to Share vision continues beyond the project and evolves into a concrete European R&I direction. Building on its cross-fertilisation activities, demonstrators and ecosystem of complementary projects, NexTArc intends to mobilise the INSIDE and Chips JU communities around a dedicated strategic track—and potentially future thematic calls or call topics—addressing the convergence of Circular Cities and Regions, Energy-Efficient Cities and Communities, Sustainable Mobility and CCAM, Community Hubs, and Collective Urban Intelligence. Such an initiative could bring together ECS technologies, trustworthy and energy-efficient AI, edge intelligence, digital twins, smart buildings and districts, the Digital Vehicle of the Future, semantic interoperability, circular value chains and secure data sharing under the common Dare to Share philosophy. Rather than creating

another isolated Smart City programme, the objective would be to establish cross-sector innovation environments in which technologies developed across Chips JU and complementary European programmes can be cross-fertilised, integrated, validated and scaled through real urban and industrial applications. NexTArc therefore aims to use its remaining activities to consolidate requirements, engage projects, cities, industry and research communities, and translate the resulting vision into concrete recommendations for future ECS SRIA priorities and Chips JU programming, turning Dare to Share from a foresight concept into a potential European.

Acknowledgement

We want to express our sincere thanks to all the researchers, industry experts, project representatives and participants who contributed to the NexTArc Cross-Fertilisation Workshop on 8th of June, 2026, Brussels. Special thanks go to the speakers and organisers: Ronald Maandonks (Signify), Maha Karim-Hosselet (INSIDE Association), Tolga Baykal (Togitek), Prof. Dr. Matthias Haase (ZHAW – Zurich University of Applied Sciences), Gianluigi Ferrari (University of Parma), Charlie Gullström (Sweco), Alper Kanak (ERARGE), Ali Serdar Atalay (AI4SEC), Ömer Öztin (AI4SEC), Salih Ergün (ERGTECH), Paolo Azzoni (Eurotech Group) and NexTArc Project Coordinator Baran Çürüklü (Mälardalen University).

Alper Kanak, Ph.D.
NexTArc Technical Coordinator / R&D Director
Ergünler R&D Co. (ERARGE)

Alper Kanak, Ph.D., is a computer scientist, R&D director and programme manager with over 25 years of experience in European research and innovation. His expertise spans trustworthy AI, cybersecurity, cyber-physical systems, digital twins, semantic technologies, smart cities and sustainable mobility. He has led, coordinated or contributed to 50+ EU-funded initiatives, including recently NexTArc, OPEVA, AI4SWEng, ESCALATE and VALU3S, bridging strategic research with industrial implementation.

Charlie Gullström
Dissemination, Communication and Exploitation Workpackage Leader [BC1.1]/ Head of Research and Innovation, Adjunct Professor in Circular Economy for Architecture and Planning
SWECO / Chalmers Technical University

Following an extensive academic career over twenty years, Gullström leads Sweco's externally funded research and innovation to support cities in urban and digital transition. She is an experienced research leader of quadruple helix initiatives to accelerate sustainable urban development and digital innovation. In NexTArc she leads dissemination activities (WP9) and coordinates Sweco's efforts as a use case host of an innovation arena in the Stockholm city district Tidningskvarteren.se.

Baran Çürüklü
NexTArc Project Coordinator / Senior Lecturer
Mälardalen University

Baran Çürüklü is a senior lecturer in artificial intelligence at Mälardalen University, Sweden. His research spans artificial intelligence, robotics, human–robot interaction, multi-agent systems, and autonomous cyber-physical systems. He has led and contributed to European research projects funded through Horizon 2020, ECSEL/Chips JU, and the European Defence Fund. His work focuses on intelligent coordination, mission planning, responsible AI, and autonomous systems, with strong emphasis on collaboration with industrial partners *across* Europe.

Matthias Haase
Scientific and technical dissemination responsible [BC1.1]/ Professor
Zurich University of Applied Sciences (ZHAW)

Prof. Dr Matthias Haase is an experienced engineer with many realized sustainable façade design and BIPB projects, coordinator of several EU projects (FP7, H2020, Horizon Europe) with degrees in Mechanical Engineering (BSc and Dipl.-Ing. from England and Germany), sustainable resource management (MEng from Germany) and Architecture (PhD from Hong Kong).
Since Sep 2020 Matthias is Professor for Building Systems and Sustainable Resource Management at the Institute of Facility Management (IFM).

Research Project Highlight

Making industrial AI visible and impactful for Europe's manufacturing future

How AIMS5.0 supports Europe's transition toward smarter, more resilient and sustainable industry

Anna Lackner
Founder & CEO, Nexuswelt

Europe's manufacturing industry is facing a period of profound transformation. Artificial intelligence, advanced semiconductor technologies and digital manufacturing systems are increasingly shaping industrial production, competitiveness and technological resilience across Europe. At the same time, European industry must respond to global competition, supply chain disruptions, sustainability targets and the need for stronger digital sovereignty. In this context, AIMS5.0 supports the transition toward more intelligent, connected and sustainable manufacturing ecosystems.

AIMS5.0: bringing AI into real manufacturing environments

AIMS5.0, "Artificial Intelligence in Manufacturing leading to Sustainability and Industry 5.0," is a major European innovation action funded under the Chips Joint Undertaking and coordinated by Infineon Technologies AG. Bringing together 53 partners and 20 industrial use cases across Europe, the project focuses on integrating advanced AI technologies into real manufacturing environments while supporting Europe's Industry 5.0 vision.

What makes AIMS5.0 particularly relevant is its strong industrial focus. The project is not limited to theoretical AI research, but demonstrates how artificial intelligence can create measurable impact within real production environments across

semiconductor and manufacturing value chains.

Strengthening Europe's semiconductor and ECS ecosystem

AIMS5.0 is embedded in a wider European innovation ecosystem shaped by organisations such as INSIDE Industry Association, AENEAS, EPOSS and the Chips Joint Undertaking. Together, these actors strengthen Europe's Electronic Components and Systems community, support industrial collaboration and help connect research results with strategic priorities in semiconductors, digital technologies and advanced manufacturing. Through this ecosystem, AIMS5.0 contributes to stronger industrial capabilities, resilient value chains and more competitive manufacturing systems.

From technical innovation to practical industrial impact

The project addresses several key European priorities, including improved production efficiency, more reliable quality control, more sustainable manufacturing processes and increased industrial resilience. A strong example is the collaboration between KAI and Infineon on AI-powered defect image classification for semiconductor manufacturing. In semiconductor production, visual inspection is a highly sensitive task: small deviations on wafers, components or production-related image data can indicate process variations, material issues or quality risks. Traditional inspection workflows can be time-consuming and depend strongly on expert interpretation, especially when large volumes of image data must be reviewed consistently.

Within AIMS5.0, this work explores how AI can support the classification of defect images by learning from structured and annotated industrial image datasets. The approach combines image preprocessing, feature extraction and machine-learning-based classification to identify relevant defect patterns more efficiently. Instead of only detecting that an anomaly exists, the system can support the categorisation of defect types and provide more consistent information for quality assessment.

From a technical perspective, such AI-based inspection workflows depend on high-quality training data, clear labelling of defect classes, validation against production-relevant image sets and continuous evaluation of classification accuracy. Explainability is also important: visual outputs, confidence values or highlighted image regions can help engineers understand why the system suggests a certain classification. This makes the AI result more useful for industrial experts and supports trust in the technology.

Cross-project collaboration with Arrowhead fPVN and SC4EU

AIMS5.0 has also demonstrated the value of cross-project collaboration within the wider European innovation ecosystem. Through activities such as the Deep Tech Collaboration Workshop in Lisbon, the project engaged with initiatives including Arrowhead fPVN and SC4EU to exchange expertise on interoperability, AI integration, digital engineering and future industrial architectures.

These collaborations encourage alignment between complementary technological

initiatives. Arrowhead fPVN contributes perspectives on flexible production value networks and interoperable industrial architectures, while SC4EU adds relevance to Europe's semiconductor value chain and skills ecosystem.

The role of INSIDE IA, AENEAS, EPOSS and Chips JU

European organisations such as INSIDE Industry Association, AENEAS, EPOSS and the Chips Joint Undertaking play an important role in this context. They create the environment in which ambitious collaborative projects can connect, exchange knowledge and contribute to Europe's technological agenda. For AIMS5.0, this ecosystem connection is essential because the project's impact goes beyond individual use cases. It contributes to a broader European effort to connect industrial AI, semiconductor innovation, digital engineering and manufacturing transformation under a shared strategic framework.

Supporting Europe's competitiveness and strategic autonomy

The impact of AIMS5.0 extends beyond individual technological results. By supporting AI-enabled manufacturing solutions, strengthening industrial collaboration and contributing to Europe's semiconductor ecosystem, the project helps reinforce Europe's long-term competitiveness and strategic autonomy in critical industrial domains. This matters in a global environment where advanced manufacturing and semiconductor technologies are increasingly linked to economic resilience, innovation leadership and technological independence.

Looking ahead: toward a more connected European industrial future

As Europe continues investing in digital transformation, semiconductors and sustainable industry, initiatives like AIMS5.0 demonstrate how collaborative innovation can support more resilient, intelligent and human-centric manufacturing systems. The project contributes to advancing industrial AI technologies and to shaping a stronger, more connected European industrial future.

As AIMS5.0 moves toward its final project phase, the General Assembly in Madrid on 16–17 June has provided an important moment to review key exploitable results, present selected use cases and strengthen collaboration with related European initiatives. The meeting also supported the transition from project implementation toward broader exploitation, visibility and long-term industrial impact across Europe's manufacturing and semiconductor ecosystem.

Relevant links

- AIMS5.0 Official Website: https://aims50.eu/
- AIMS5.0 News & Stories: https://aims50.eu/index.php/news-stories
- AIMS5.0 at INSIDE Connect 2025: https://aims50.eu/index.php/news-stories?view=article&id=228&catid=12
- Deep Tech Collaboration Workshop: https://aims50.eu/index.php/news-stories?view=article&id=234&catid=12

The next cybersecurity race has already started

Memory-safe computing and Europe's opportunity to lead the next Secure-by-Design computing platform

Mike Eftimakis
Founding Director
CHERI Alliance

As cyber threats accelerate and software complexity grows, Europe faces more than a security challenge: it faces a major economic opportunity. By leading the transition from reactive cybersecurity to computing platforms that prevent vulnerabilities by design, Europe can strengthen its digital resilience. It can also create new markets, attract investment, and build high-value skills and supply chains. In doing so, Europe can turn its research leadership in memory-safe computing into lasting industrial advantage.

Europe's digital future is increasingly being shaped by forces that extend far beyond technology itself. Geopolitical competition, cyber conflict and artificial intelligence are transforming software and computing infrastructure into strategic assets. Security can no longer be viewed as a purely technical concern; it has become a matter of economic resilience, industrial competitiveness and technological sovereignty.

Geopolitical tensions

Today, geopolitical competition increasingly manifests itself through technology.

A century ago, industrial power depended primarily on access to raw materials, manufacturing capacity and energy resources. In the twenty-first century, semiconductors, software platforms and digital infrastructure have joined that list of strategic assets. Nations that control critical technologies increasingly enjoy economic, military and diplomatic advantages.

The semiconductor industry provides perhaps the clearest illustration. Processor architectures, fabrication technologies and software ecosystems determine who develops the products that power communications, transportation, critical infrastructure, artificial intelligence and defence systems.

As a result, governments are investing unprecedented resources into securing technology supply chains, developing domestic capabilities and reducing strategic dependencies.

Europe has responded through major initiatives, including the European Chips Act and an expanding cybersecurity regulatory framework. Yet sovereignty involves far more than manufacturing capacity. It depends on controlling the technologies that underpin future digital systems.

The cyber battlefield

At the same time, the threat landscape is evolving faster than ever.

For many years, cybersecurity was viewed primarily as a corporate issue. Today it has become a matter of national resilience. Cyber operations now accompany military conflicts, diplomatic tensions and economic competition. The distinction between warfare, espionage and criminal activity is increasingly blurred.

Attacks no longer target only governments or military organisations. Energy operators, hospitals, transportation providers, industrial facilities and telecommunications networks can become targets because disruption itself has strategic value.

Every increase in digitalisation therefore has two consequences: an increased **productivity**, but also an increased **dependency**. The same technologies that improve efficiency can also create new vulnerabilities.

Cybercrime has also become industrialised. Vulnerability discovery, exploit development and ransomware operations are now highly organised businesses operating at global scale. The time between vulnerability disclosure and exploitation continues to shrink.

Meanwhile, the attack surface keeps expanding. Electricity grids, transportation networks, industrial facilities, vehicles,

medical systems and communications infrastructure all increasingly depend on software. As digitalisation accelerates, more systems become critical infrastructure.

The AI acceleration effect

Artificial intelligence may prove to be one of the largest amplifiers of both opportunity and risk.

Modern AI systems rely on vast software stacks involving operating systems. A single AI workload may involve millions of lines of code spanning operating systems, hypervisors, networking layers, model frameworks, inference engines and hardware accelerators. Every additional layer creates new opportunities for vulnerabilities.

AI systems themselves are becoming attractive targets. As organisations increasingly rely on AI for automation and decision making, compromising these systems can have significant consequences.

At the same time, AI is becoming a force multiplier for attackers. Vulnerability discovery, code analysis and exploit development, tasks that once required teams of highly skilled researchers, can increasingly be automated and scaled. This raises a growing concern: attackers find vulnerabilities faster than defenders can fix them.

The trend has become sufficiently concerning that a coalition of cybersecurity leaders behind the Zero Day Clock initiative[1] warns that the time between vulnerability disclosure and active exploitation could continue shrinking dramatically. Based on observed trends in thousands of exploited vulnerabilities, the initiative projects that median time-to-exploitation may fall below one hour and potentially below one minute during 2027, driven in part by increasingly capable AI-assisted vulnerability analysis and exploit development. The project has attracted support from recognised figures across the cybersecurity community, including security practitioners, CISOs, researchers and industry pioneers.

Ironically, much of the discussion around trustworthy AI focuses on ethics, transparency and governance while overlooking a fundamental question: can AI systems be trusted if the underlying computing infrastructure remains vulnerable?

Trustworthy AI ultimately requires trustworthy computing foundations.

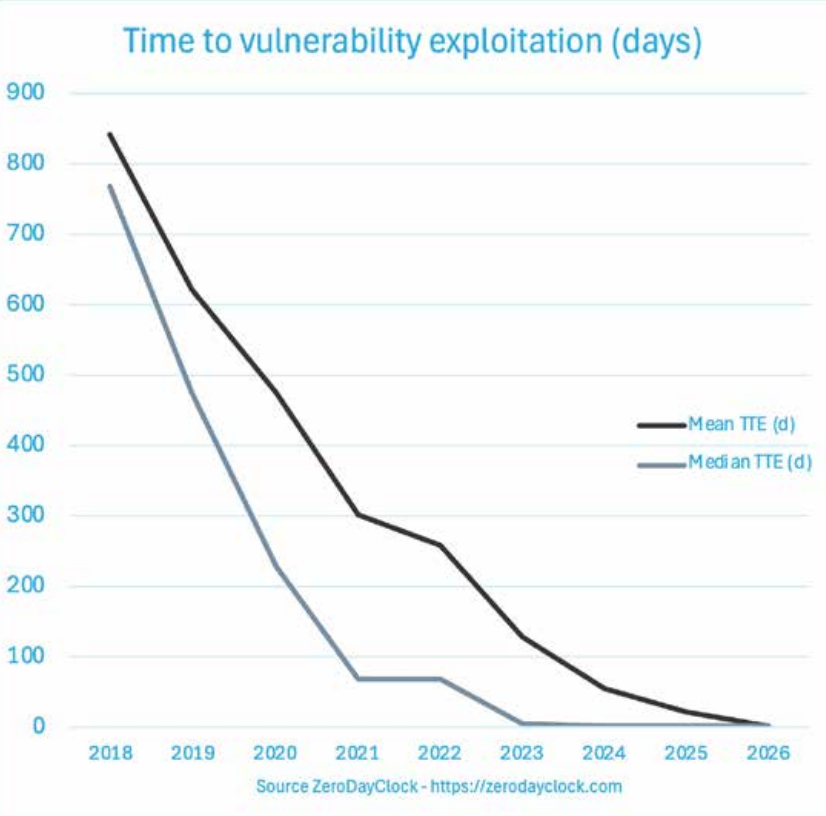

The challenge beneath the challenge

Together, geopolitical competition, industrialised cybercrime, expanding attack surfaces and accelerating software complexity create a perfect storm for digital security.

Europe's resilience increasingly depends on the security of its computing foundations. Yet despite decades of investment and effort, the cybersecurity problem remains stubbornly unsolved.

Why cybersecurity is still failing

The vulnerability paradox

Over the past three decades, cybersecurity has become a major global industry. Organisations have deployed more tools, more monitoring systems, more compliance frameworks, more regulations and more security processes than ever before. Security spending continues to increase year after year.

Yet vulnerabilities continue to accumulate[2]. Security teams face growing backlogs. Software suppliers struggle to patch systems quickly enough. The number of reported vulnerabilities continues to rise.

The industry has become increasingly effective at detecting problems, while remaining remarkably ineffective at preventing them.

Legacy software and supply chain risk

A major part of today's cybersecurity challenge lies in the software foundations of modern society. Critical systems depend on software that has evolved over decades, including operating systems, industrial controllers, networking equipment, automotive platforms and telecommunications infrastructure. Collectively, these systems contain trillions of lines of code that have been developed, modified and maintained over many years. Replacing them wholesale is economically unrealistic. Instead, organisations must continue securing software that was never designed for today's threat landscape and threat actors.

Much of this software is written in C and C++, languages that remain essential because they provide direct hardware control and exceptional performance. However, they also require developers to manage memory explicitly, a task that becomes extraordinarily difficult when systems reach millions of lines of code and involve thousands of developers over long periods of time. Buffer overflows, use-after-free vulnerabilities and related memory-safety errors continue to appear even in highly scrutinised software. This is not primarily a failure of developers; it is an inevitable consequence of managing complexity at enormous scale.

That complexity is further amplified by modern software supply chains. Today,

When one vulnerability reaches millions

Modern software depends on shared libraries, packages and toolchains. That reuse accelerates development, but it also creates concentration risk: a flaw or compromise in one low-level component can spread across thousands of downstream products at once.

SolarWinds (2020) demonstrated the power of this approach. Attackers compromised the build environment of a widely used network management product and inserted malicious code into legitimate software updates. When customers installed those trusted updates, the compromise spread to numerous government agencies and major enterprises.

Log4Shell (2021) highlighted a different form of supply-chain risk. A critical vulnerability in the widely used Log4j logging library affected hundreds of thousands of applications and services around the world. Many organisations were unaware that their products even depended on the vulnerable component.

XZ Utils (2024) revealed yet another attack vector. A carefully crafted backdoor was introduced into a widely used open-source compression library after a long-term effort to gain the trust of the project's maintainers. Although the issue was discovered before widespread deployment, it demonstrated how attackers can target the social and governance aspects of open-source ecosystems.

These incidents share a common lesson: software supply chains provide attackers with leverage and scale. By compromising a single trusted component, update mechanism or development process, attackers can potentially reach thousands—or even millions—of downstream systems.

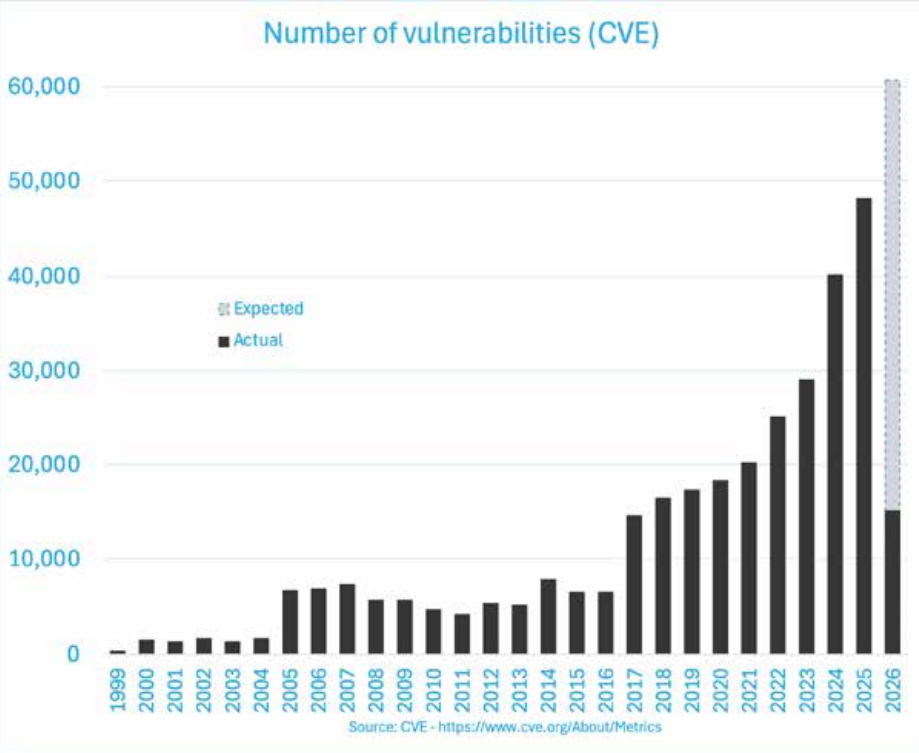

almost no organisation develops software entirely from scratch[3]. Applications are assembled from operating systems, libraries, open-source components, frameworks and third-party packages that collectively form vast dependency networks. A vulnerability introduced anywhere in this ecosystem can affect countless downstream products and organisations. Attackers increasingly exploit software supply chains (cf sidebar), because they provide leverage and scale, allowing a single compromise to reach thousands of targets. The challenge is therefore no longer simply protecting individual applications, but securing entire software ecosystems whose complexity continues to grow year after year.

The memory safety problem

Most people outside software engineering assume vulnerabilities are primarily caused by mistakes in authentication, configuration or cryptography. While these issues certainly matter, one of the most persistent problems lies deeper inside the software itself.

Many analyses identify **memory-safety vulnerabilities** as the single largest source of software security flaws. Defects such as buffer overflows, use-after-free vulnerabilities and invalid memory accesses continue to affect even the most widely used and heavily reviewed software systems, accounting for more than **70%** of reported vulnerabilities according to multiple converging studies from Google, Microsoft and other organisations[4 5].

As highlighted previously, these defects often originate from the complexity of modern software. The vulnerabilities may remain invisible for years, but once discovered, attackers can deliberately exploit them to bypass security controls, execute malicious code or gain unauthorised access to systems.

What makes these vulnerabilities particularly problematic is their longevity. Many have remained exploitable despite decades of research, code reviews, security testing, static analysis tools, runtime checks and exploit mitigations.

In many cases, software components have accumulated layers of defensive mechanisms intended to make exploitation harder. These protections often come at the cost of additional complexity, memory usage and performance overhead, yet the underlying memory-safety vulnerabilities continue to reappear, forcing the industry into an endless cycle of mitigation and patching.

This raises an important question: are we building increasingly sophisticated defences around fundamentally insecure foundations?

Why the situation is getting worse

Patching has become the dominant

cybersecurity response mechanism. Keeping pace with newly discovered vulnerabilities was already challenging, but several trends are making the problem increasingly difficult. AI-assisted analysis tools are uncovering previously unknown vulnerabilities in mature software that has been deployed for years or even decades. In many cases, organisations lack complete visibility into where vulnerable components are deployed, particularly when software bills of materials (SBOMs) are incomplete, outdated or absent. Identifying affected systems, assessing risk and deploying patches at scale can therefore become a major operational challenge.

Organisations also face a practical reality: there is insufficient time, budget and skilled labour to investigate and remediate every vulnerability. Security teams are increasingly overwhelmed by patch backlogs that grow faster than they can be addressed.

The challenge is even greater in sectors with long product lifecycles. Industrial facilities, critical infrastructure, aerospace systems and vehicles often remain operational for decades, far exceeding traditional software support cycles. Continuous updates are not always technically, operationally or economically feasible. As support ends, vulnerable software often remains in service for years, exposing the organisations and critical infrastructure connected to it to ongoing risk.

As vulnerability discovery accelerates and the cost of remediation continues to rise, the industry's reliance on patching looks increasingly unsustainable. We have become highly effective at finding vulnerabilities, but far less successful at eliminating the conditions that create them in the first place. We continue to treat symptoms while the root causes remain largely unchanged.

A new generation of computing security

From detection to prevention

For decades, cybersecurity has focused primarily on detecting attacks, patching vulnerabilities and limiting the impact of successful intrusions. While these capabilities remain essential, they do little to address the underlying causes of many security flaws. A growing number of researchers, governments and industrial organisations therefore advocate a shift towards **security by design**: building security into systems from the outset rather than adding defensive layers after deployment. The objective is not simply to respond to vulnerabilities more effectively, but to prevent entire classes of vulnerabilities from being introduced in the first place. This represents a fundamental transition from reactive cybersecurity towards computing platforms that are secure by construction.

One of the most promising opportunities for security by design lies in addressing memory-safety vulnerabilities, which remain the primary source of software security flaws as highlighted previously. After more than fifty years of increasingly sophisticated software tools, testing methodologies and defensive mitigations, the industry continues to struggle with the same underlying classes of vulnerabilities. This suggests that the challenge is not simply one of better tooling, but of architectural foundations. Memory-safe computing therefore takes a different approach: rather than relying exclusively on software mechanisms to detect unsafe behaviour, it introduces protection directly into the hardware. The goal is not merely to make attacks more difficult, but to prevent entire classes of vulnerabilities by construction.

CHERI

One of the most promising solutions to the memory-safety problem is **CHERI** (Capability Hardware Enhanced RISC Instructions)[6], a computing architecture specifically designed to address the root causes of many software vulnerabilities. The fundamental insight behind CHERI is deceptively simple. Traditional processors treat memory addresses largely as numbers, leaving it to software to ensure that those addresses are used correctly. CHERI transforms memory references into protected capabilities that carry not only the location of data, but also information about how that data may be used. Each capability includes hardware-enforced bounds and permissions, ensuring that software can only access memory within explicitly authorised limits. In effect, permissions become attached directly to memory references, and the processor continuously verifies that software remains within those boundaries.

This seemingly modest architectural change enables a fundamentally different security model in which hardware actively participates in enforcing software security. Instead of relying solely on software checks, testing tools and defensive mitigations to detect errors after they occur, CHERI moves critical security guarantees into the underlying architecture itself. The result is a platform capable of mitigating large classes of memory-safety vulnerabilities by construction. Rather than merely raising the cost of attack, CHERI seeks to pre-emptively eliminate entire categories of vulnerabilities at their source, making it one of the most comprehensive and practical approaches yet developed for

How CHERI prevents out-of-bounds accesses

Traditional processors treat pointers much like ordinary numbers, making it possible for software to access memory beyond an intended data buffer and create vulnerabilities such as buffer overflows.

CHERI replaces pointers with hardware-protected capabilities, which become the only means of accessing memory. Each capability carries hardware-enforced bounds and permissions defining precisely which memory can be accessed and how. Since capabilities cannot be forged in software and every memory access is validated by hardware, memory-safety attacks can no longer be exploited – unauthorised memory accesses are prevented by construction.

improving the security of existing software systems.

Perhaps the most important aspect of CHERI is its compatibility with existing software ecosystems. Organisations cannot simply discard or rewrite vast codebases developed, validated and refined over decades. CHERI therefore offers an incremental path: much existing software can gain immediate protection through recompilation, typically with minimal changes. Stronger protection then comes gradually from modifying software to use finer-grained compartmentalisation and CHERI's native hardware security features over time.

This evolutionary approach is further supported by CHERI's hybrid mode, which allows CHERI-enabled and legacy code to execute side by side within the same system. Recompiled software immediately benefits from additional hardware-enforced security guarantees, while existing binaries continue to operate unchanged. By combining significantly improved security with a realistic and low-risk adoption model, CHERI lowers the barriers to deployment and makes large-scale industrial adoption far more achievable than approaches that require complete software rewrites.

In supply-chain attacks such as those described in the sidebar, CHERI would not necessarily prevent a compromised component, malicious update or vulnerable dependency from entering a system. Its value lies in reducing what happens next: by enforcing memory safety, least privilege and fine-grained compartmentalisation, CHERI can help limit the blast radius of a compromise and make it harder for an attacker to move from one component to another or corrupt unrelated parts of the system.

CHERI should be viewed as a necessary element of a broader secure-by-design toolbox. It combines with complementary techniques across three layers: hardware techniques such as hardware roots of trust, physical security and cryptographic accelerators; software techniques such as secure boot, trusted execution environments, compartmentalisation and runtime checks; and process or assurance techniques such as software bills of materials, formal verification, secure development practices, vulnerability disclosure processes and supply-chain assurance. Among the various approaches being explored to improve software security, CHERI stands out not only because of its technical merits, but also because of its origins: it is one of the few foundational computing innovations of recent decades in which Europe has played a leading role from the outset.

Europe's unusual advantage and the risk of losing it

Europe helped create the technology

Much of the foundational work behind CHERI originated at the University of Cambridge, where researchers collaborated with international partners to explore how hardware architecture could improve software security. Over the past fifteen years, that early academic work has grown into a broad community spanning academia, industry and government. Teams across Europe and the United Kingdom continue to play leading roles in advancing the technology, while a growing number of organisations have contributed processor implementations, operating-system support, compiler toolchains, formal verification technologies and open-source software.

European companies such as Codasip (Germany)[7] and SCI Semiconductor (United Kingdom)[8] have helped demonstrate how CHERI concepts can be translated into commercial products and deployable technologies. At the same time, researchers across the United Kingdom, Belgium, Germany, Switzerland, the Netherlands, France and other European countries continue to expand the expertise surrounding memory-safe computing. Europe has also become a hub for open-source CHERI development, with projects such as lowRISC's CHERIoT Ibex (United Kingdom)[9] and the OpenHW Foundation's CVA6-CHERI[10] helping turn the technology into practical, accessible processor platforms.

Public investment has also played a critical role in moving CHERI beyond the laboratory. Government-funded programmes in the UK[11] have supported experimental hardware platforms, software infrastructure, standards activities and real-world demonstrations, helping to bridge the gap between academic innovation and industrial deployment. As a result, Europe's role extends across the full value chain, from architecture to software and tools, open-source communities and the skills base required to support adoption. Few emerging computing technologies can claim such a comprehensive European footprint, making memory-safe computing a rare area in which Europe begins from a position of genuine technological leadership.

Complementary strengths

Europe possesses several advantages that naturally align with the emergence of secure-by-design computing. The continent maintains world-class expertise across semiconductors, embedded systems and processor design, with globally recognised strengths in chip architecture, semiconductor design, safety-critical software and systems engineering. Europe also plays a leading role in industries such as automotive, aerospace, rail and industrial automation, sectors where safety, reliability and certification are fundamental requirements rather than optional features. These industries have decades of experience designing systems to prevent failures before they occur and may therefore be particularly receptive to architectural approaches that seek to eliminate classes of cybersecurity vulnerabilities by design.

Europe also possesses an asset that is often perceived as a weakness: a strong regulatory framework. As cybersecurity and resilience increasingly become mandatory requirements through initiatives such as the Cyber Resilience Act and evolving critical infrastructure regulations, compliance may become a powerful market driver for secure-by-design technologies. In this context, Europe's industrial strengths, safety culture and regulatory environment could combine to create favourable conditions for the adoption of memory-safe computing at scale.

The danger of invention without industrial leadership

Technological leadership and economic leadership are not the same thing. Europe has repeatedly produced world-class research and breakthrough innovations, but invention alone does not secure long-term economic or strategic advantage. Value is captured when research becomes products, supply chains, standards and markets. That transition is already beginning around CHERI, with companies and organisations in the USA, Japan, China and other countries exploring or integrating it into future processors, platforms and products. Europe's window to turn early leadership into industrial strength will therefore not remain open indefinitely. Scientific excellence creates the opening; deployment determines who captures the jobs, investment, expertise and economic value.

This distinction matters especially for memory-safe computing. The technology emerged from years of pioneering research, but its long-term impact will depend on the

ecosystem that forms around it. Successful platforms need more than technical innovation: they need skilled engineers, mature software tools, semiconductor products, integrators, suppliers, training programmes, standards bodies and customers ready to deploy at scale. Ecosystems, not individual breakthroughs, determine whether a technology becomes a market.

History suggests that speed matters as much as excellence. Once a technology moves from research into commercial use, advantages can compound rapidly: early users build expertise, product launches attract customers and partners, universities create pools of trained engineers, and first-wave standards often shape entire industries. Momentum becomes increasingly difficult to reverse.

The key question is therefore no longer who invented memory-safe computing, but who will industrialise it. Who will deploy it first in commercially significant products? Who will build the processor IP, development tools and software base required for adoption? Who will train the next generation of engineers, establish supply chains, certification processes and standards, and become the preferred partner for organisations seeking secure-by-design computing platforms?

The next cybersecurity race is therefore not primarily scientific. The concepts have been demonstrated, the research base is mature, and deployment is beginning. **The real competition now is industrial**. The winners will be those who build the products, skills, markets and supporting ecosystems around the technology. For Europe, the stakes are high: it can lead a new generation of secure computing, but it could also see others capture the economic and strategic benefits of a technology it helped create.

Europe's opportunity: from research to impact

A new security advantage for European industry

Security is rapidly becoming a product requirement rather than a technical feature. Automotive manufacturers, defence contractors, telecom operators, AI platform providers and operators of critical infrastructure increasingly need stronger assurances about the software that underpins their products and services. In these sectors, cybersecurity failures can have operational, financial and even safety consequences.

Secure-by-design computing therefore has the potential to become a powerful competitive differentiator, much as functional safety became a defining characteristic of many European industries over recent decades. Organisations that can demonstrate stronger security properties by construction rather than through layers of defensive software may gain a significant advantage in markets where trust, reliability and resilience are increasingly valued.

This transition could also create opportunities to **reshape existing market positions**. Major technology shifts often favour new entrants because they require new products, new expertise and new ways of engineering systems. Secure-by-design computing requires new processor architectures, optimised software stacks, development tools and verification methodologies. As a result, the transition towards memory-safe computing could create opportunities for European semiconductor companies, software providers and system integrators to establish leadership in emerging markets. At a time when Europe is seeking to strengthen its role in semiconductor design and capture a greater share of future value creation, secure computing may represent an opportunity not merely to participate in the next generation of digital infrastructure, but to help define it.

Turning European Constraints into Strategic Advantages

Europe's regulatory environment is often portrayed as a burden on innovation. Yet history shows that regulation can also create markets. Safety requirements helped transform automotive engineering, environmental regulations stimulated advances in energy efficiency, and quality standards helped create globally competitive industrial sectors. Cybersecurity and digital resilience requirements may have a similar effect.

As initiatives such as the Cyber Resilience Act and related regulations increasingly require stronger security assurances, technologies capable of delivering those assurances gain a natural market advantage. Secure-by-design computing therefore aligns closely with Europe's broader policy objectives. Rather than viewing cybersecurity compliance as a cost, Europe has an opportunity to use regulatory frameworks to stimulate demand for technologies that address vulnerabilities at their source. In this scenario, regulation becomes not merely a mechanism for reducing risk, but a catalyst for industrial innovation, investment and market creation.

What Europe should do next

To capture the value of CHERI, Europe must shift from research to deployment:

- Launch deployment-oriented pilots
- Back security-by-design in commercial chips
- Fund projects to develop tools, skills and migrate software
- Prove value in industrial, telecoms and infrastructure use cases
- Align standards, procurement and regulation around secure-by-design computing

This opportunity is amplified by the fact that memory safety sits beneath many of the technological transitions currently reshaping society. Artificial intelligence depends on trustworthy computing infrastructure. Post-quantum cryptography requires secure software implementations. Software-defined vehicles, industrial automation systems, critical infrastructure and defence modernisation programmes all rely on increasingly complex software platforms, where security supports safety requirements. Memory-safe computing does not compete with these trends; it strengthens them. As a foundational technology, it has the potential to improve the security and resilience of an entire generation of digital systems.

Building the Ecosystem

Realising this opportunity, however, requires more than technological success. The transition to secure-by-design computing cannot be achieved by researchers alone. It depends upon collaboration across an entire ecosystem that includes semiconductor companies, software providers, tool developers, system integrators, regulators, open-source communities, universities and end users. Europe's challenge is therefore not simply technological; it is organisational.

Today, many security requirements originate from operators of critical infrastructure, industrial companies and end users, while research priorities often emerge

from academic programmes or public funding initiatives. Europe needs stronger mechanisms to connect these worlds. The most successful technology ecosystems create continuous feedback loops that translate operational needs into research priorities and transform breakthroughs into deployable products. Without these connections, even strong technical leadership can struggle to generate lasting industrial impact.

This is where ecosystem organisations can play a crucial role. Organisations such as the **CHERI Alliance**[12], a non-profit industry association dedicated to accelerating CHERI adoption, help bring together stakeholders from across the value chain, creating forums for collaboration, knowledge sharing, standards development and industrial adoption. Recognising the importance of connecting secure-by-design computing with Europe's wider innovation and industrial communities, the CHERI Alliance recently joined the **INSIDE** ecosystem. The objective is not merely to promote a technology, but to strengthen the connections between academics, semiconductor companies, software developers, policymakers, system integrators and end users required to translate technological advances into large-scale industrial impact. Ultimately, the challenge is no longer proving that memory-safe computing works; it is ensuring that the expertise, products, standards and supply chains needed to support widespread adoption emerge quickly enough for Europe to convert technological leadership into lasting economic and strategic advantage.

Conclusion

Memory-safe computing represents more than another cybersecurity technology. It reflects a broader shift in how digital systems are designed, moving security from an ongoing process of detection, patching and mitigation towards architectures that prevent entire classes of vulnerabilities by construction. As software continues to underpin artificial intelligence, critical infrastructure, transportation, communications and defence systems, the security of computing foundations is becoming a strategic concern rather than a purely technical one.

Europe enters this transition from an unusually strong position. The continent helped create many of the technologies underpinning modern memory-safe computing, possesses world-class expertise in semiconductors, safety-critical engineering and cybersecurity, and benefits from a growing ecosystem of industrial, academic and public-sector stakeholders. At the same time, evolving policy initiatives aimed at strengthening cybersecurity and resilience are increasing demand for secure-by-design approaches across multiple sectors.

Yet technological leadership does not automatically translate into industrial leadership. The decisive phase is no longer invention, but industrialisation. The organisations, regions and ecosystems that successfully transform memory-safe computing into products, skills, standards and industrial capabilities will shape the next generation of digital infrastructure.

History has shown that major technology transitions create rare opportunities to redefine competitive positions. Memory-safe computing may be one of those moments. As cybersecurity becomes a strategic concern across artificial intelligence, critical infrastructure, defence, telecommunications and industrial systems, the ability to deliver security by construction could become a powerful competitive advantage.

The next cybersecurity race has therefore already begun. It is no longer a race to identify vulnerabilities faster, nor even a race to invent new defensive technologies. Increasingly, it is a race to eliminate vulnerabilities by design and to industrialise the platforms capable of doing so. Europe has contributed significantly to creating this opportunity. The challenge now is to ensure that it also captures the economic, technological and strategic benefits that follow.

Key insights

- Cybersecurity is increasingly a matter of resilience, competitiveness and sovereignty
- Memory-safety vulnerabilities remain one of the most persistent sources of software risk
- Secure-by-design computing aims to eliminate vulnerabilities rather than continually mitigate them
- CHERI provides a practical path towards memory-safe computing through hardware-enforced protection and incremental adoption
- Europe helped create the technology; the next challenge is industrialising it at scale

1 Available: zerodayclock.com.

2 "CVE Metrics," Available: https://www.cve.org/About/Metrics.

3 K. e. a. Reid, "Towards Sustainable and Secure Reuse in Dependency Supply Chains," 2025. Available: https://arxiv.org/html/2503.02804v3.

4 Microsoft, "Trends, challenge, and shifts in software vulnerability mitigation," Available: https://github.com/microsoft/MSRC-Security-Research/blob/a0b9de6d9feef9d6b518afbb-311baa108efdd169/presentations/2019_02_BlueHatIL/2019_01%20-%20BlueHatIL%20-%20Trends%2C%20challenge%2C%20and%20shifts%20in%20software%20vulnerability%20mitigation.pdf.

5 CISA, "The Urgent Need for Memory Safety in Software Products," 2023. Available: https://www.cisa.gov/news-events/news/urgent-need-memory-safety-software-products.

6 P. G. N. J. W. M. R. H. A. J. A. e. a. R. N. M. Watson, "Capability Hardware Enhanced RISC Instructions: CHERI Instruction-Set Architecture (Version 9)," Available: https://www.cl.cam.ac.uk/techreports/UCAM-CL-TR-987.pdf.

7 Codasip. Available: https://codasip.com.

8 SCI Semiconductor, Available: https://www.scisemi.com.

9 lowRISC, "Ibex," Available: https://lowrisc.org/ibex/.

10 "CVA6-CHERI," Available: https://github.com/Capabilities-Limited/cheri-cva6.

11 UKRI, "Digital security by design," Available: https://www.ukri.org/what-we-do/browse-our-areas-of-investment-and-support/digital-security-by-design/.

12 CHERI Alliance, Available: https://cheri-alliance.org.

Making Europe's Cloud-Edge-IoT ecosystem easier to see, use and reuse

Veronica Vuotto
Communication, Dissemination & Outreach Specialist

Almost two years after its launch in October 2024, the Horizon Europe project CEI-Sphere has entered a different phase of work. The first part of the project was largely about understanding Europe's Cloud-Edge-IoT (CEI) landscape: analysing markets, mapping stakeholders and building a shared vocabulary through resources such as the Market Briefs and the Hourglass Model. With the project now heading towards its conclusion in March 2027, the priority has shifted towards making that knowledge visible, practical and reusable by others.
The last few months show this shift clearly. A first version of the use cases catalogue went online, a new liaison with the European logistics community was established, the first CEI-Sphere Hackathon took place in Cagliari, and two webinars opened up market and security discussions to a wider audience. In parallel, a new Horizon Europe project, INCLUDE, has started to carry part of this work forward.

Seeing the ecosystem: the Use Case Catalogue

In May, CEI-Sphere launched its Use Case Catalogue, an interactive tool that gives a clear overview of the use cases being developed within the two Large-Scale Pilots, O-CEI and COP-PILOT. The idea behind is that. Europe's CEI ecosystem is growing quickly, but it is not always easy to see how the different pieces fit together. The catalogue brings together use cases from energy, mobility, logistics, industry and agrifood, and allows users to explore them by sector, geographical location, stakeholder group or technology maturity level. Next to each use case, the tool also shows the organisations involved, from technology providers and research centres to industrial and public actors.

More than a list, the catalogue works as an entry point into the ecosystem. It makes visible who is working with whom, which technologies recur across sectors, and where interoperability and standardisation actually make a difference. It also arrives at a moment when the European CEI market, estimated at around 158 billion euros in 2024, keeps growing while moving solutions beyond the pilot stage remains difficult.

Connecting with logistics: the ALICE Liaison Programme

In July, CEI-Sphere joined the ALICE Liaison Programme, the network of the European Technology Platform for Logistics. The aim is to create a stable channel between the CEI community and the people who run and design logistics operations in Europe. Logistics depends more and more on real-time data, distributed infrastructures and automated decision-making, which means that the ability of different systems to work together is no longer a technical detail. CEI technologies can help, by allowing data to be processed and exchanged closer to where it is generated and needed. CEI-Sphere and ALICE share several priorities, including interoperability and standardisation, the digitalisation of supply chains, trusted data sharing and Data Spaces, automation and intelligent transport systems.

The liaison connects naturally with CEI-Sphere's Mobility, Logistics and Software-Defined Vehicles Sphere, which gathers use cases and stakeholders working on freight transport, connected mobility and real-time logistics. In practice, the cooperation will involve sharing results and use cases, taking part in thematic activities and discussing how open digital infrastructures relate to ALICE's vision of the Physical Internet. Tools such as the Use Case Catalogue and the Hourglass Model give the conversation something concrete to work with.

Code the Continuum: the first CEI-Sphere Hackathon

On 30 June and 1 July, the University of Cagliari hosted Code the Continuum, the first CEI-Sphere Hackathon, organised together with O-CEI, COP-PILOT and the Eclipse Foundation. For two days, students, developers, researchers and open-source contributors worked side by side on real challenges involving the movement of data and the orchestration of software across cloud, edge and IoT environments.

Three tracks were on offer. The first asked teams to orchestrate distributed applications with OpenSlice in a simulated cloud-edge environment. The second explored how Eclipse Zenoh can support energy-aware edge-cloud applications, comparing cloud-only and edge-accelerated deployments. The third, proposed by the Eclipse Foundation, invited teams to build a Jakarta EE application using AI and agentic tools to model enterprise domains through the Hourglass Model.

The two days closed with fourteen final pitches and a total of 4,500 euros awarded across the three tracks. Some teams may also gain further visibility through CeADAR, Ireland's Centre for Applied AI, or explore with the French space technology company Parsimoni how their solutions could be adapted for satellite and in-orbit applications. Beyond the prizes, the Hackathon was a useful test of something the project talks about often: technologies developed inside the Large-Scale Pilots were picked up by people from outside them and turned into working prototypes in two days.

Two webinars: adoption and trust

July also brought two webinars. The first, on 8 July, was dedicated to accelerating the adoption of Cloud-Edge-IoT ecosystems. Organised with the BRIDGE initiative and drawing on CEI-Sphere's Deep Dive Analysis, it looked at how CEI solutions are actually being adopted in the energy, mobility and industry spheres, what prevents them from scaling beyond pilots, and how open source, open standards and the CEI Trust Label can help. The session also connected these findings with the work of the BRIDGE Business Model Working Group, bringing business sustainability and commercialisation pathways into a discussion that often stays technical.

The second webinar, on 23 July, addressed new challenges in security and privacy for CEI. Speakers presented the work of the INSTAR project on cybersecurity and digital identity standards, discussed Data Spaces as governance frameworks rather than pure technology, and explained why post-quantum cryptography is becoming a practical concern for distributed infrastructures. Pilot 7 of O-CEI showed how a secure Vehicle-to-Grid charging ecosystem can be built over private 5G networks, combining AI-based intrusion detection, dynamic trust management and zero-trust principles. The recurring message was that security, privacy and interoperability have to evolve together rather than one after the other.

Where to meet CEI-Sphere next

After the summer, the project will be present at two events. The first is the third edition of AIOTI Days, taking place in Florence on 28 and 29 September 2026, with a programme focused this year on generative AI and interoperable intelligent systems. Straight after that, CEI-Sphere will join INSIDE Connect 2026 in Palermo, on 30 September and 1 October.

A new chapter: the INCLUDE project

With CEI-Sphere ending in March 2027, the question of continuity is a fair one. Part of the answer is the INCLUDE project, which has recently kicked off under Horizon Europe. INCLUDE acts as the coordination backbone of the EURO-3C Initiative on Connected Collaborative Computing Networks, aligning demand-side and supply-side stakeholders across the Telco-Edge-Cloud continuum.

There is a direct connection with our community: INCLUDE builds on the results, methodologies and network developed over the years by EU Cloud Edge IoT and more recently by CEI-Sphere and its Large-Scale Pilots, and several familiar faces from the ecosystem are part of the new consortium. Its planned outputs will look recognisable too, for instance, a stakeholder technology roadmap, a shared governance framework, a catalogue of collaborative tools and open-source integration mechanisms, a technical handbook and training activities. In the coming months, expect joint conversations between the two projects on roadmaps, open-source building blocks and vertical engagement.

For the remaining part of its lifetime, CEI-Sphere will keep supporting O-CEI and COP-PILOT and promoting their open calls, expanding the Use Case Catalogue as new use cases mature, and working with standardisation communities, open-source foundations and sectoral platforms such as ALICE. The ambition is to leave behind a set of living resources, a catalogue, a shared model, a community and a working relationship with the projects that come next, that other people can pick up and keep using.

Online version is available at Inside-association.eu

Publisher
INSIDE Industry Association
High Tech Campus 69-3
5656 AG Eindhoven, The Netherlands

Design and Creative lay-out
Studio Kraft – Veldhoven, the Netherlands

Acknowledgements
With thanks to the interviewees, project participants, INSIDE Industry Association office, the INSIDE Industry Association Presidium and other INSIDE Industry Association-involved persons for any assistance and material provided in the production of this issue of the INSIDE Magazine.

Contributions
The INSIDE Industry Association office is interested in receiving news or events in the field of Intelligent Digital Systems. Please submit your information to info@Inside-association.eu

INSIDE
Industry Association

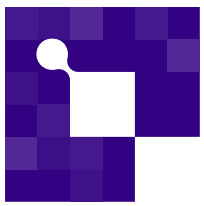

INSIDE
Industry Association

INSIDE-association.eu