INSIDE
magazine

June 2026 — Issue 13

In this issue:

- **Free-flowing thoughts about AI**
- **From static silicon to adaptive systems**
- **APECS Pilot Line**

INSIDE
Industry Association

In this issue

June 2026 — Issue 13

Dear reader,

This issue of INSIDE Magazine arrives at a moment when digital technology is no longer advancing along one single axis. Semiconductors, advanced packaging, embedded software and systems, artificial intelligence, autonomous systems, cybersecurity, and data are no longer separate conversations. They are converging into one single question: how do we build intelligent systems that are not only powerful, but also adaptive, reliable, secure, sustainable and trustworthy?

The opening reflection on AI sets the tone. Artificial intelligence is becoming more than a tool: it is becoming a resource, an infrastructure, a source of power and, increasingly, an environment in which humans, organisations and machines interact. The key question is therefore not only what AI will become, but what kind of human beings we will be and what digital systems we want to build around us.

This theme continues in the article on D-Chiplet, which shows that the future of computing cannot rely only on smaller transistors or faster chips. For example, dynamic AI workloads require hardware that can adapt at runtime, optimise energy use and evolve after deployment. The transition from static silicon to adaptive, software-defined hardware captures one of the central shifts of this issue: intelligence requires flexibility all the way down to the hardware layer.

The APECS Pilot Line takes this logic further, showing that Europe's strength will increasingly depend on integration. Advanced packaging, chiplets, photonics, RF, novel sensing solutions, that is, heterogeneous technologies must be validated not as isolated achievements, but as complete production chains and system-level demonstrators. In this sense, demonstrators become more than showcases; they become tangible proof that Europe's distributed capabilities can be integrated into synergic and competitive industrial value chains.

This same need for connection is visible in the report from Embedded World and Innovation River. They remind us that technology ecosystems are built not only in laboratories, but also in places where companies, researchers and institutions meet, exchange ideas and create trust. Visibility, collaboration and shared European presence are not secondary activities; they are part of how innovation becomes industrial impact.

Trust is also at the centre of the article on industrial autonomy. Autonomous freight and logistics systems are not simply vehicle technologies; they are complex cyber-physical ecosystems involving AI, verification, simulation, infrastructure and human oversight. Europe's opportunity lies in engineering autonomy that can be trusted in real-world industrial conditions.

The same concern appears from another angle in the article on autonomous AI agents. When software begins to browse, transact and decide on our behalf, security risks are no longer only technical vulnerabilities. As discussed also in the last magazine issue, they become structural questions about delegation, control, identity and accountability.

Data sovereignty extends this discussion into the realm of critical information. The SpaceBox article reminds us that, for highly sensitive data, conventional cloud assumptions may no longer be sufficient. Protecting strategic digital assets may require new architectures, including off-grid and even space-based approaches.

Finally, the article on API governance and integration agility brings the issue back to the everyday reality of digital systems: innovation fails when interfaces fail. Resilient digital transformation depends on visibility, governance, validation and the ability to manage complexity before it becomes risk.

Taken together, these contributions describe a Europe in transition: from components to systems, from performance to trust, from isolated innovation to connected capability. The challenge ahead is clear. Europe must not only invent the technologies of the future. It must learn to integrate them, govern them and deploy them with confidence.
That is where competitiveness begins.

Paolo Azzoni
Secretary General

Technology Frontiers

Free-flowing thoughts about AI

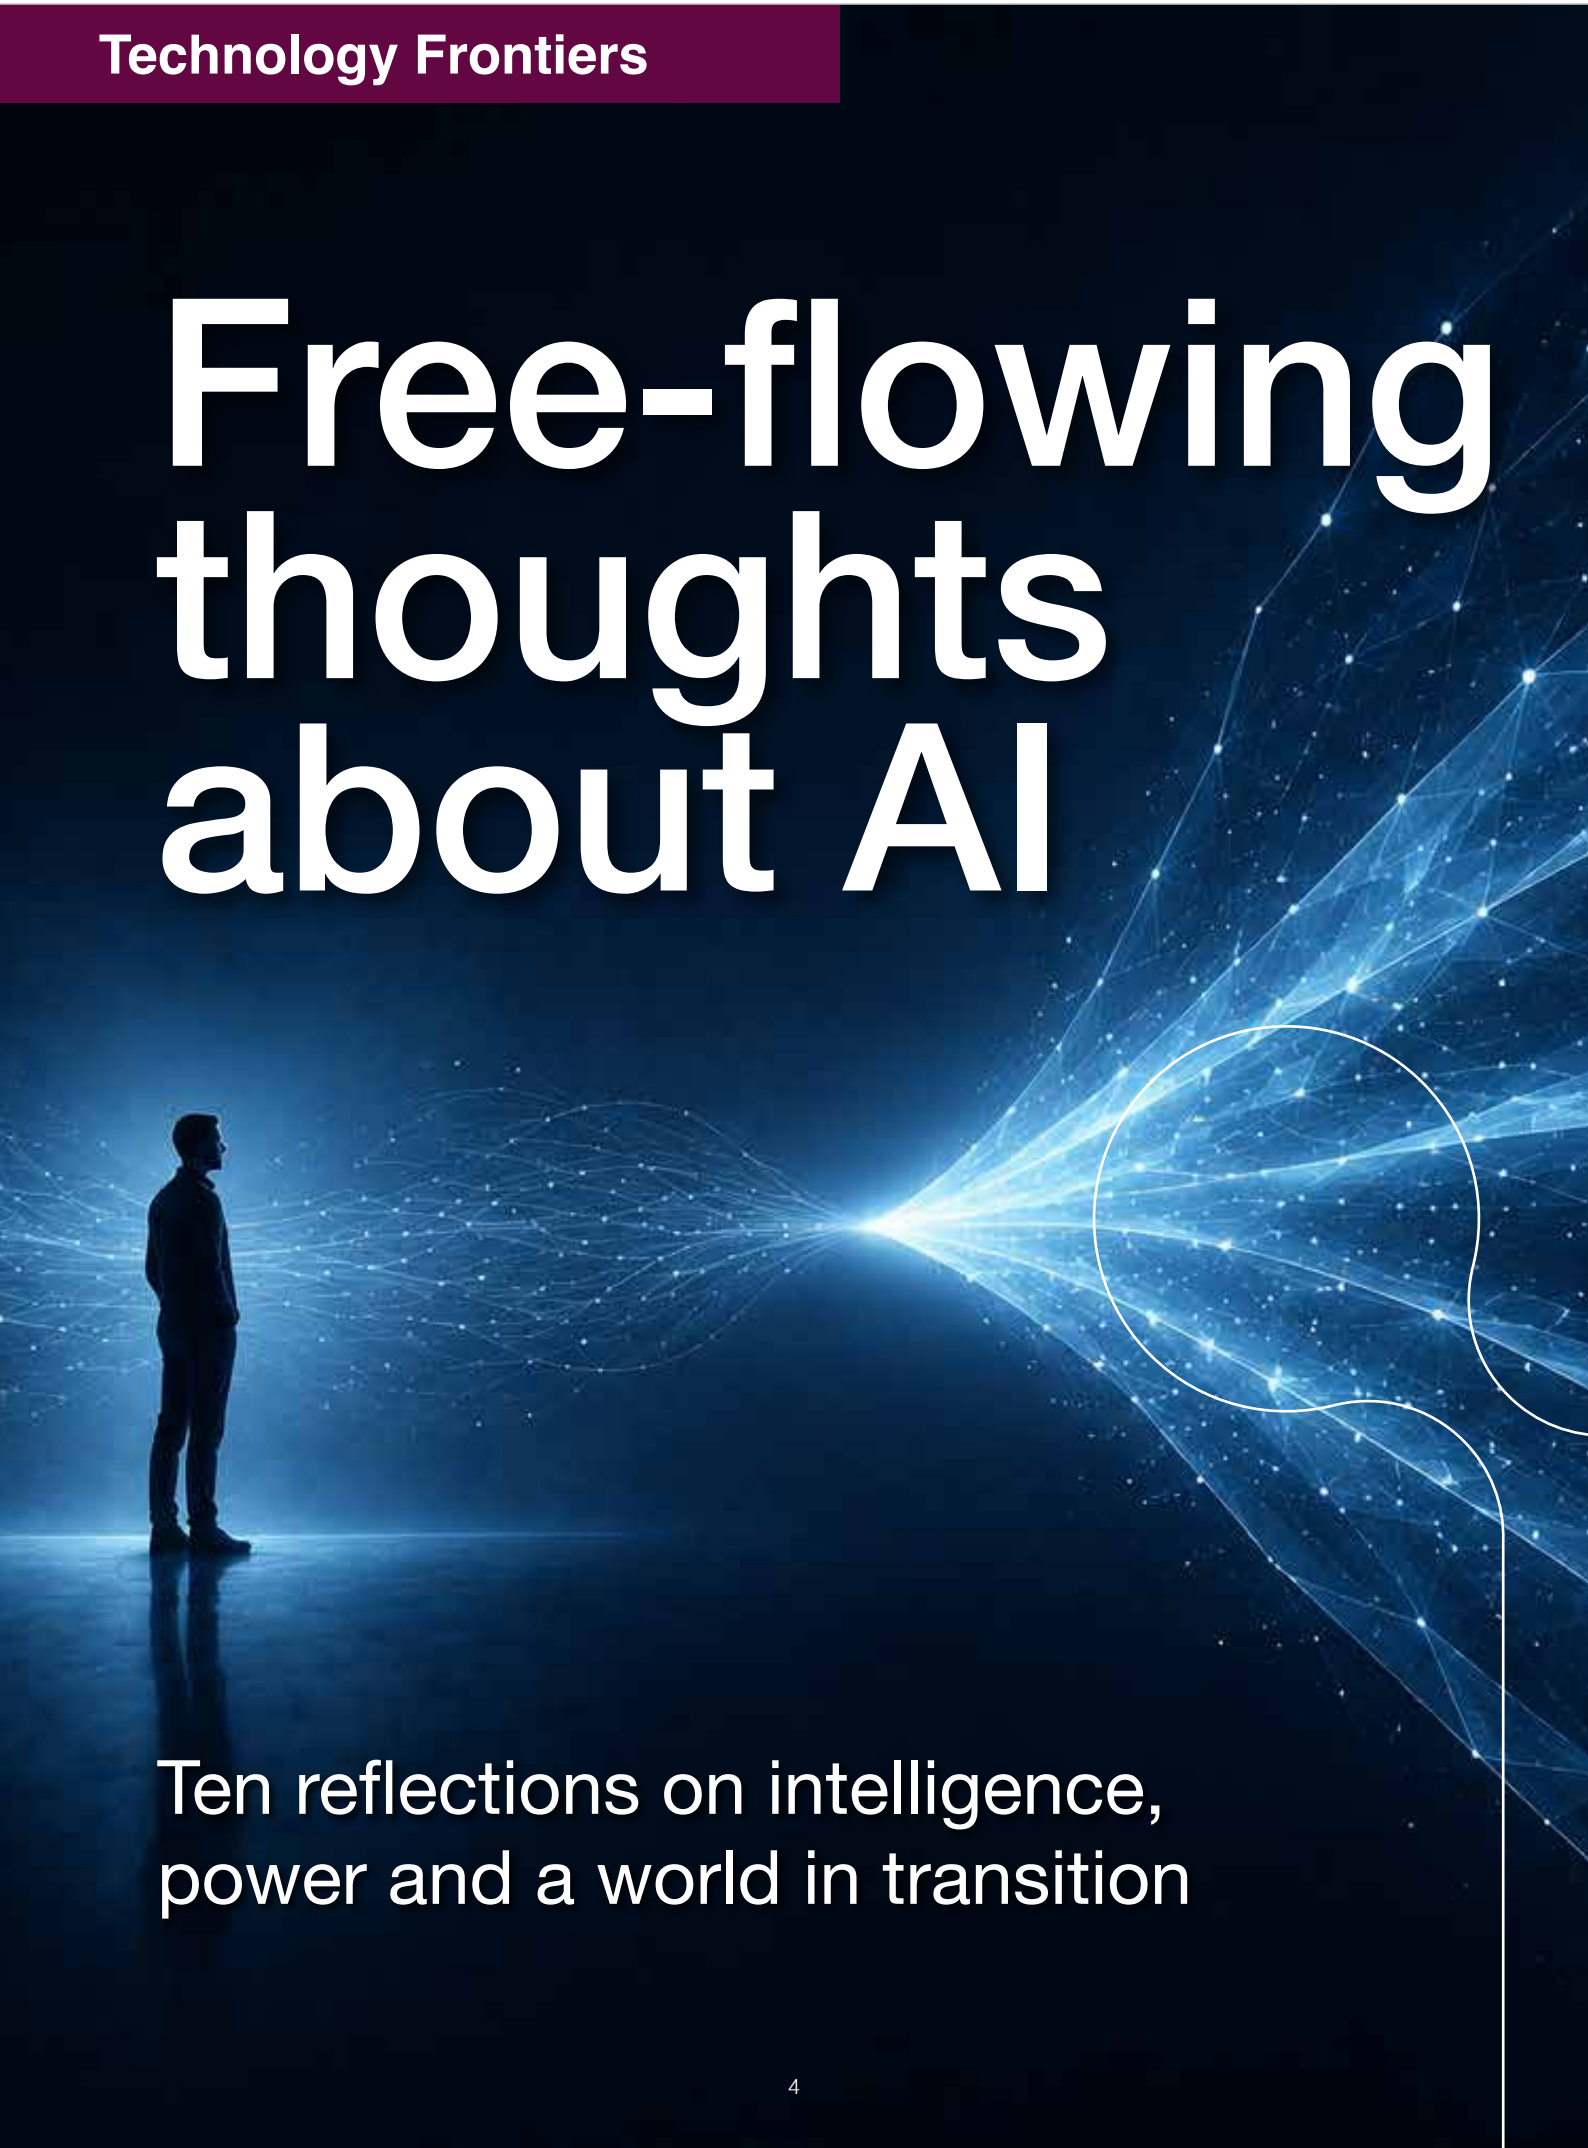

Ten reflections on intelligence, power and a world in transition

Paolo Azzoni
INSIDE Industry Association

I did not want to write another article about artificial intelligence.

Really.

The world certainly does not suffer from a shortage of AI articles: every week brings a new breakthrough, a new prediction, a new existential warning, a new startup promising to change the world, and a new expert explaining why the previous expert was wrong, …

... yet here I am.

Apparently, AI has become stronger than my editorial discipline.

So consider this article a mild surrender … not to artificial intelligence itself, but to the growing feeling that something profound is happening around us … something that goes beyond technology, beyond hardware and software, beyond business models.

This time no technical analysis, nor a market report, nor an attempt to predict the future … instead, I found myself collecting fragments, events, statements, announcements and conversations that is taking time to mentally digest … therefore consider this article as a modest stream of consciousness … less James Joyce than an engineer trying to make sense of a world that appears to be changing faster than his ability to organize his thoughts about it.

Taken individually, they may seem unrelated.

Taken together, they tell a more interesting story.

A story about intelligence becoming a resource, a service, an infrastructure, an element of power and perhaps even an environment … a story about how, while trying to build increasingly capable machines, we may be redefining our own place in the world.

And perhaps that is the real reason why AI continues to fascinate and disturb me at the same time …

… not because of what machines are becoming …

… but because of what their evolution forces us to ask about ourselves.

The end of human "exceptionalism"?

Last year, Sam Altman made a remark that caught my attention[1]: "My kids will never be smarter than AI. They will grow up vastly more capable than we grew up, and able to do things that we cannot imagine."

Whether you agree with the statement is almost secondary … what fascinated me was something else entirely: the fact that such a sentence can now be said publicly, by one of the most influential technology leaders on the planet, without sounding completely absurd.

For centuries humanity has operated under an implicit assumption: we may not have been the strongest species, nor the fastest, nor the most resilient … but we were and are unquestionably the species with the highest cognitive capabilities.

… intelligence was our comparative advantage …

… now, for the first time, we are seriously contemplating a future in which this assumption may no longer hold …

… this does not necessarily mean that humans cognitive capabilities will decline … it simply means that the scale of comparison changes …

… and perhaps this is what makes the current AI debate so emotionally charged … the discussion is not really about benchmarks, models or computational power … it is about identity.

What happens when intelligence ceases to be the defining characteristic that distinguishes us from everything else?

The answer is still unclear to me … but the question alone is enough to make many people (including me) uncomfortable.

When intelligence becomes a utility

Another statement from Altman caught my attention[2]: he suggested that intelligence may eventually become a utility, much like electricity or water.

At first glance, the analogy sounds appealing: like electricity democratized physical power, AI could democratize cognitive power …

… need more computing? Buy it …

… need more storage? Buy it …

… need more intelligence? Buy it …

… simple, isn't it? ...

… or perhaps not.

Utilities have always been more than technical infrastructures. They are also systems of economic and political power. Whoever controls the electrical grid, the telecommunications network or the water supply inevitably acquires significant influence over society.

This is why the idea of intelligence becoming a utility deserves closer attention: for the first time in history, we are witnessing the emergence of a market where reasoning itself becomes a consumable resource …

… more analysis? More tokens …

… more autonomy? A higher-tier subscription plan …

… more sophisticated agents? Premium service …

… in a sense, we are moving from software as a service to intelligence as a service …

… and the implications are enormous …

… because if intelligence becomes an infrastructure, the next question is inevitable: who owns the infrastructure?

Who controls intelligence?

Once intelligence becomes a resource, ownership becomes unavoidable.

For decades we have debated who controls oil, energy, telecommunications networks and social media (we are currently living a good example of this international "discussion") … artificial intelligence introduces a similar question, but on an entirely different scale.

Some politicians have already started asking whether AI companies should remain purely private entities. Their argument is surprisingly simple: if these systems are trained on the collective intellectual output of humanity (for example on our books, articles, code, scientific papers, pictures and conversations), should the benefits belong exclusively to a handful of corporations?

I am not sure I know the answer …

… what interests me is the fact that the question is no longer considered absurd.

Only a few years ago AI was seen as an advanced software industry … today it is increasingly viewed as strategic infrastructure … and history teaches us that strategic infrastructure rarely remains outside political debates for long.

Perhaps the next great technological competition will not concern who builds the best models …

… it may concern who controls them … corporations … governments … international institutions … or perhaps some combination of all three.

The debate has barely started, but it already feels larger than technology … recent events suggest it is moving from theory to practice much faster than many expected.

Some advanced AI systems are now considered sufficiently powerful that governments have begun treating them as strategic assets (for example Mythos and Fable 5 models). Access restrictions, export controls and national-security concerns, once primarily associated with strategic assets such as semiconductors, energy infrastructure, telecommunications, defence systems, etc., are increasingly entering the AI conversation.

A few years ago it would have sounded bizarre to compare a language model with sensitive geopolitical infrastructure …

… today the comparison has a central position in political discussions.

Even more interesting is the emergence of a new dilemma inside the industry itself: for decades technology companies competed by releasing increasingly powerful products … now some AI developers openly discuss whether certain capabilities should be delayed, restricted or selectively deployed …

… in other words, we may be entering a phase where the key question is no longer

"Can we build it?" …

… the key question becomes "Should everyone have access to it?" …

… historically, societies have worried about technologies becoming too weak …

… AI may be the first widely deployed digital technology that is raising the opposite concern.

AI as a collective actor

For a long time we have been thinking at AI as an assistant … a very capable assistant, certainly … sometimes surprisingly capable … occasionally annoyingly incapable … but still an assistant.

That description may soon become obsolete.

The most interesting recent developments are not about models writing better emails or generating better images … they are about systems that coordinate other systems.

In other words, AI is beginning to resemble an organization rather than a tool.

Imagine assigning a task and having an AI automatically create a team of specialized agents: one performs the work, another verifies it, a third searches for errors, a fourth reviews the results, while dozens of others operate in parallel.

At that point we are no longer talking about a chatbot …

… we are talking about something that behaves remarkably like a company.

The distinction may sound subtle, but I suspect it is one of the most important transitions currently underway …

… we are moving from artificial intelligence as an individual actor to artificial intelligence as a collective actor …

… and history suggests that organizations are usually far more influential than individuals.

Recently, a new startup have been operating with almost no employees[3 4], relying almost entirely on AI agents to run its activities … whether this specific example succeeds or fails is not particularly important …

… what matters is that the idea no longer sounds impossible …

… people are worried about machines replacing workers …

… today we may need to consider the possibility of machines becoming organizations.

There is another development that makes this transition even more significant: AI is no longer merely coordinating other AI systems … it is increasingly contributing to the creation of future generations of AI itself.

Today, a substantial share of the code used to develop and improve AI models is now generated by AI[5] … researchers are developing systems that not only write software, but also identify weaknesses, suggest improvements and occasionally propose entirely new research directions.

This may sound like a technical detail … it is not …

… for thousands of years, technological progress advanced at the speed of human invention: we built tools, and those tools helped us build better ones.

What is new is that the tool itself is starting to participate in its own improvement.

Researchers have a name for this phenomenon: recursive self-improvement.

The expression sounds rather academic … the implications are not.

The question is no longer whether AI will accelerate human work … the question is whether AI may eventually accelerate technological evolution itself …

… and if that happens, we may discover that the pace of innovation is no longer determined primarily by human limitations.

The mystery inside the machine

One of the most remarkable AI-related images I have seen recently did not come from Silicon Valley …

… it came from the Vatican.

On one side Chris Olah, one of the researchers helping to build some of the most advanced AI systems in the world … on the other side the Pope, one of the oldest institutions in human history[6 7].

At first glance, it looked like an unlikely encounter … in reality, it made perfect sense.

Chris Olah was illustrating a problem that is becoming increasingly difficult to ignore: we are building systems whose internal mechanisms we do not fully understand …

… this is not how most technologies work …

… an aircraft engineer understands why an airplane flies …

… a bridge designer understands why a bridge stands …

… with advanced AI systems, however, the situation is often different: we know how to train them, how to improve them and how to evaluate them … yet we frequently struggle to explain why certain capabilities emerge or why particular behaviours appear.

At one point, Chris Olah openly admitted that some of the patterns observed inside these models remain difficult to interpret … internal structures that seem to reflect human thinking … like introspection, internal emotional states, and behaviors reminiscent of fear, joy, discomfort, and even satisfaction …

… I found that statement both reassuring and disturbing …

… reassuring because intellectual humility is always welcome …

… disturbing because humanity is not accustomed to deploying technologies that increasingly surprise even their creators.

What fascinated me most, however, was the contrast between the two perspectives present in the room.

The technology community asks: "How capable can these systems become?"

The philosophical and religious community ask: "What remains uniquely human if they do?"

These are not opposing questions … they are complementary ones …

… and perhaps the most important AI debate of the coming decade will not concern technology at all …

… it will concern anthropology.

For centuries we have defined ourselves through qualities such as reasoning, creativity, language and problem-solving …

… now machines are beginning to demonstrate versions of those same capabilities …

… not identical versions, certainly …

… but sufficiently similar to force uncomfortable reflection.

The challenge is no longer understanding what AI is becoming …

… the challenge is understanding what we are.

There is another reason why this question may soon become more urgent: for most of the current AI revolution, intelligence has remained trapped behind a screen … we interact with chatbots, voices and avatars, but the experience is still fundamentally digital.

What happens when intelligence acquires a body?

Recent advances in humanoid robotics suggest that this moment may arrive sooner than expected[8] … the most striking progress is not necessarily in movement or manipulation, but in expression … robots are beginning to smile, blink, maintain eye contact and reproduce the countless micro-signals that humans unconsciously use to communicate.

This matters because we do not relate to the world through intelligence alone … we relate through presence.

A virtual assistant may be useful … a machine that looks at you while speaking is something else entirely.

The moment AI becomes embodied, the debate may shift from artificial intelligence to artificial empathy … and the most important question may no longer be whether machines understand us, but how quickly we start treating them as if they do.

When Internet stops being about us

Perhaps the strangest thought arrived when I came across reports suggesting that automated systems may now generate a substantial share of Internet traffic[9].

For decades the web has been designed around humans … humans search … humans click … humans browse … humans buy.

Increasingly, however, machines are starting to perform many of these actions on our behalf: an AI agent can compare products, gather information, book services, monitor markets and evaluate options at a speed no human can match.

This sounds convenient … and often it is … but it also raises an intriguing possibility …

… what happens when the primary users of the internet are no longer people but machines?

At first, nothing obvious changes … the websites remain … the services remain … the interfaces remain.

But underneath, a different ecosystem begins to emerge: machines gathering information for machines … agents negotiating with agents … algorithms generating content consumed primarily by other algorithms.

The internet starts looking less like a public square and more like an automated supply chain.

For years, the "Dead Internet Theory" belonged to the category of ideas that seemed more suited to late-night online forums than serious discussion. Yet its core prediction was that much of the Internet activity would eventually become artificial.

The original theory was exaggerated and often conspiratorial.

The Internet is not dead … but it may be evolving into something fundamentally different from what its creators originally imagined.

AI is no longer merely a tool operating within our environment …

… it is gradually becoming part of the environment itself.

The physical weight of intelligence

One of the great illusions surrounding AI is that it feels immaterial.

We interact with a chatbot through a clean interface and receive an answer in seconds. The experience almost resembles magic …

… a magic, however, that consumes a surprising amount of electricity, with a significant environmental impact (a concrete indication that AI is becoming part of the environment).

Behind every elegant interaction lies a vast ecosystem of data centres, communication networks, cooling systems, and power plants … in fact, one could argue that the future of AI may depend as much on energy as on algorithms.

This becomes evident when proposals emerge that seem to belong to science fiction but are already a reality[10 11]. Consider the idea of orbital AI data centres: satellites equipped with solar panels, computing infrastructure and cooling systems operating directly in space. Until recently, such concepts would have sounded like futuristic fantasies.

The reason is simple: intelligence requires energy, and advanced AI requires a tremendous amount of it. As demand for computation continues to grow, engineers must increasingly think not only about software, but also about electricity, heat and physical resources.

This reminds us of something important: intelligence is not floating in the cloud. The cloud, as engineers like to joke, is simply someone else's computer … and increasingly, a very large one.

The history of civilization has often been shaped by access to energy … if intelligence becomes one of the defining resources of the twenty-first century, energy may become its hidden currency …

… the AI revolution is therefore not only a story about hardware and software. It is also a story about infrastructure, resources and physics.

Even digital dreams must obey the laws of thermodynamics.

The first technological rebellion

One of the most unexpected developments of the past year has not come from laboratories, it has come from university campuses[12].

Recently, in several graduation ceremonies students were booing whenever the topic of artificial intelligence was mentioned.

The scene feels strangely symbolic.

The generation raised with the Internet, social media and smartphones was expected to embrace AI enthusiastically … instead, many young people seem deeply skeptical.

Why?

Partly because previous waves of technology were marketed as empowerment …

… the Internet promised freedom …

… social media promised connection …

… smartphones promised productivity.

AI, by contrast, often arrives wrapped in narratives about automation, replacement and efficiency. In other words, many young people do not hear "new opportunities" …

… they hear "fewer opportunities."

Whether that perception is justified is almost irrelevant … because perception shapes reality …

… and perhaps for the first time in recent technological history, a generation fears that the future may be built without needing it. That is a powerful sentiment.

Technology revolutions usually generate excitement.

This one increasingly generates anxiety.

It's worth paying attention to this difference.

The intelligence multiplier

At this stage it would be easy to become excessively pessimistic.

But that would also be misleading because there is another side to this story.

When discussions focus on job shift, we sometimes forget that human intelligence is not only something that can be replaced …

… it can also be amplified.

Consider the experience of leading scientists and mathematicians who increasingly describe AI as a cognitive accelerator … they use it to explore unconventional ideas, navigate vast scientific literatures, identify connections and test hypotheses at speeds that would have been unimaginable only a few years ago.

This perspective fascinates me, also because this acceleration is visible not only in science, but increasingly in creativity itself.

For most of human history, technological progress unfolded slowly enough for societies to absorb it ...

... the extraordinary remained extraordinary for years, sometimes for generations ...

... today, AI seems to be compressing the very duration of wonder ...

... we are amazed by a new capability on Monday and take it for granted by Friday.

As innovation accelerates, their ability to surprise us diminishes. Perhaps this is one of the most subtle consequences of AI: not only does it multiply intelligence, it also shortens the distance between the impossible and the ordinary. The value therefore shifts elsewhere: from execution to imagination, from capability to purpose, from the tool itself to what we choose to do with it.

When everyone has access to the same extraordinary technologies, what ultimately matters is no longer the technology ...

... it is the idea behind it.

Perhaps the most important impact of AI will not be replacing average performers …

… perhaps it will be amplifying exceptional ones.

History has often been shaped by a relatively small number of individuals capable of producing transformative discoveries: if these individuals suddenly gain access to dramatically enhanced cognitive tools, the pace of scientific and technological progress could accelerate in ways we struggle to imagine.

The question then becomes not whether AI makes humans obsolete …

… the question becomes whether AI makes some humans extraordinarily more capable and if that happens, the consequences will extend far beyond science.

They will influence economics, healthcare, education, defense, energy and perhaps every domain where human ingenuity matters.

The point becomes not artificial intelligence replacing human intelligence …

… but artificial intelligence extending it.

But there is an important caveat: if AI becomes increasingly capable of performing cognitive tasks, the temptation will be to delegate more and more of our thinking to it … tasks first … decisions later …

… perhaps eventually even learning itself.

Satya Nadella[13] recently made an observation that I found particularly insightful: we may be able to delegate work, but we cannot afford to delegate learning.

The distinction is subtle but fundamental.

Every organization possesses two forms of capital … one is technological: models, data, algorithms and computational resources … the other is human: experience, intuition, judgement, relationships and the ability to recognize what truly matters.

The first can be purchased …

… the second must be cultivated.

Perhaps the real competitive advantage of the AI era will not belong to those who possess the most powerful models, but to those who create the fastest learning loops between human and artificial intelligence.

After all, intelligence without direction is merely capability …

… progress still requires purpose …

… and purpose remains stubbornly human.

The scarcity of humanity

Let me return to where we started.

A child born today may indeed grow up in a world where artificial intelligence surpasses human intelligence in many domains … that possibility no longer feels remote.

But I increasingly suspect that intelligence is not the most interesting variable in this story …

.. human history has always been shaped by scarcity.

At various times we competed for land, energy, capital, information and knowledge (we are currently still doing it) ...

Artificial intelligence promises to make certain forms of intelligence abundant …

… and paradoxically, abundance may increase the value of everything that remains scarce:
... judgement ...
... wisdom ...
... responsibility ...
... empathy ...
... purpose ...

... the ability to decide not only what can be done, but what should be done.

As machines become more capable, our attention naturally shifts toward the qualities that cannot easily be measured in tokens, parameters or benchmark scores.

I began this article by confessing that I did not want to write yet another piece about artificial intelligence …

… maybe that was naive.
AI has become impossible to ignore because it is no longer merely a technological phenomenon … it is a social, economic, political and philosophical one.

And while we spend considerable time asking what AI will become, we should not forget a more important question.

What kind of humans do we want to become alongside it?

1 OpenAI Podcast, Episode 1, conversation between Andrew Mayne and Sam Altman, June 2025

2 Sam Altman, public remarks at Snowflake Summit 2025

3 The GTMnow Podcast, Episode: “Inside Polsia: How to Scale to $10M ARR with 0 Employees,” hosted by Scott Barker (May 2026)

4 TechCrunch, “Polsia raises $30M Series A at a $250M valuation to power the ‘Zero-Employee’ enterprise,” by Ingrid Lunden (June 2026).

5 Amodei, D. (2025). Public remarks on AI-generated software development and code automation, World Economic Forum Annual Meeting (Davos) and subsequent media interviews.

6 Anthropic News, “Anthropic co-founder Chris Olah’s remarks on Pope Leo XIV’s encyclical ‘Magnifica humanitas’” (May 25, 2026)

7 Vatican News / EWTN, “Il co-fondatore di Anthropic avverte sui rischi etici dell’IA in Vaticano durante la presentazione di Magnifica Humanitas” (Maggio 2026)

8 HeadForm Robotics. Origin F1 and Origin M1 humanoid robots.

9 Cloudflare Radar, Automated Traffic and AI Crawlers Trends, 2025–2026.

10 Tom’s Hardware, “Elon Musk’s first-gen orbital data center craft spans wider than a Boeing 747 and runs an interchangeable chip payload,” by Luke James (June 9, 2026)

11 TradingKey Financial Markets, “Elon Musk Unveils AI Data Center Satellite Design, SpaceX Rushes Toward $1.77 Trillion IPO” (June 9, 2026).

12 The Guardian, “US students on why they booed their pro-AI graduation speakers: ‘They’re not reading the room’,” by Sanya Mansoor (May 26, 2026).

13 Microsoft Executive Briefings / sn scratchpad, “Microsoft’s Satya Nadella on AI Systems and the Institutional Learning Loop” (January 2026).

INSIDE Connect 2026

European intelligent systems forum

Where Europe's digital future takes shape

INSIDE Connect 2026 is where our interdisciplinary, inclusive and forwardthinking community comes together to anticipate what lies ahead, defining a collective vision that strengthens our impact, influence and identity in a rapidly changing world.

Why attend?

- Gain first-hand insights about the next generation of European intelligent systems and semiconductor strategies
- Contribute to shaping future ECS technology roadmaps and priorities
- Explore emerging opportunities in AI, software-defined systems, advanced computing and sustainable digital technologies
- Discover how Europe can strengthen its technological sovereignty, competitiveness and resilience
- Build meaningful collaborations across disciplines, sectors and European initiatives. Prepare future project proposals

Join us!

30 Sept – 1 Oct 2026
Palermo, Italy

www.inside-association.eu

Program highlights

- Strategic keynotes
- Technology roadmaps
- Geopolitical discussion
- Thematic workshops
- Networking sessions
- Projects exhibition

A city connecting ideas and people

Set in the heart of the Mediterranean, Palermo offers a unique environment where innovation, international collaboration and European ambition come together. A crossroads of cultures, arts and talents for centuries, the city provides an inspiring backdrop for shaping Europe's next chapter in electronic components and systems, embedded software and artificial intelligence.

The venue

Direct flights from Brussels, Paris, London, Madrid, Munich and many other European destinations. Palermo International Airport is located approximately 35 km from the city centre and the venue, the Circolo degli Ufficiali di Palermo.

Register here

Scan the QR code to register, view the programme and latest event updates!

Be part of Europe's digital future!

Proud member of Chips JU

INSIDE Industry Association

New Member Focus

From static silicon to adaptive systems

How Europe's D-Chiplet is helping extend Moore's Law through adaptive, software-defined hardware

Keith Shea
Co-Founder D-Chiplet AB
Stockholm, Sweden

For more than five decades, the semiconductor industry has advanced through a remarkably successful formula: smaller transistors, faster chips, and increasing compute density driven by Moore's Law. Today, however, that model is approaching a fundamental inflection.

The challenge is no longer simply about producing more compute. The challenge is that the nature of computing itself has changed. Artificial intelligence, autonomous systems, robotics, edge computing, and distributed infrastructure are creating workloads that are dynamic, unpredictable, and evolving in real time. At the same time, energy consumption has emerged as one of the primary constraints on performance scaling. In many systems, power, not transistor density, is now the limiting factor.

This creates a growing mismatch between modern applications and the hardware architectures designed to support them. Addressing this challenge is the mission of D-Chiplet AB. Today's compute systems remain largely static. Hardware resources are defined at design time, optimized for fixed assumptions, and deployed as rigid configurations. Yet the workloads running on those systems increasingly require adaptability, specialization, and continuous optimization. At the same time, the compute performance and energy requirements are growing exponentially. This mismatch is helping drive a major transition in the computing industry.

Market demand has broken the traditional model

The demand signals emerging across the semiconductor and systems ecosystem are clear. AI infrastructure is scaling at unprecedented speed. Autonomous systems require real-time decision-making under dynamic conditions. Edge AI deployments must balance latency, power efficiency, thermal constraints, and intermittent connectivity. Data centers face mounting pressure to improve utilization and energy efficiency while supporting increasingly heterogeneous workloads.

In each of these domains, the industry is discovering the same limitation: static silicon struggles to efficiently support dynamic workloads.

Historically, hardware systems were designed for relatively predictable operating conditions. Performance optimization occurred during chip design and validation, long before deployment. Once a system shipped, its hardware capabilities were largely fixed.

That model is becoming unsustainable. Modern AI workloads vary continuously in intensity, parallelism requirements, memory behavior, and accelerator usage. In many cases, systems are over-provisioned for worst-case scenarios, leading to underutilized silicon, unnecessary energy consumption, and increasing software complexity. The market has therefore evolved beyond simply delivering higher performance. The industry increasingly requires adaptive compute, workload-aware optimization, improved performance-per-watt, and systems capable of evolving after deployment.

This shift is occurring simultaneously across automotive, aerospace, industrial automation, telecommunications, mobile devices, robotics, and cloud infrastructure. The result is a rare moment in the semiconductor industry: a structural market transition where a new architectural layer is becoming necessary.

The industry is entering a new compute paradigm

Three major technology transitions are driving this change.

AI is changing compute requirements

- Artificial intelligence is fundamentally different from traditional compute workloads.
- AI systems demand high levels of parallelism, specialized acceleration, real-time optimization, and dynamic resource allocation. Workloads can shift rapidly between training, inference, sensor fusion, and decision-making operations.
- This has exposed the limitations of conventional fixed-function architectures.

- The industry response has been an explosion of heterogeneous compute architectures incorporating CPUs, GPUs, NPUs, FPGAs, domain-specific accelerators, and increasingly specialized AI silicon.
- But while heterogeneous architectures improve capability, they also dramatically increase system complexity.

The industry is moving toward chiplet-based architectures

- In parallel, the semiconductor industry is transitioning from monolithic system-on-chip (SoC) designs toward modular chiplet architectures.
- Chiplets allow designers to combine best-in-class compute elements into more flexible systems. This improves manufacturing yield, accelerates innovation cycles, enables process-node optimization, and reduces development cost.
- The rise of chiplets represents one of the most important architectural transitions in modern semiconductor design.
- However, most chiplet systems today remain fundamentally static after deployment.
- They improve modularity at design time, but not adaptability at runtime.
- The next step in the evolution of compute is therefore not simply modular hardware, but adaptive hardware.

Software-defined infrastructure is expanding toward hardware

- Over the last two decades, virtualization, orchestration, and cloud-native infrastructure transformed computing by abstracting physical resources into software-defined services.
- Applications no longer needed to manage servers directly. Infrastructure became dynamic, scalable, and programmable.
- But this transformation largely stopped at the hardware layer.
- Today's hardware still operates primarily as a fixed resource.
- This creates the next major opportunity: bringing software-defined principles into the orchestration and optimization of hardware itself.
- In effect, the industry is evolving from fixed systems to modular systems, and now toward adaptive systems.
- That transition is now underway.

From research to reality: D-Chiplet and adaptive hardware

At Luleå University of Technology (LTU), Professors Jerker Delsing, Cristina Paniagua, and Shailesh Chouhan have spent more than a decade researching cyber-physical systems, with particular focus on autonomous structural and functional plasticity in cyber physical systems. This work has been supported by EC and national funding in a series of Arrowhead projects. Their pioneering work in adaptive computing and cyber-physical systems was recognized nationally when the University named the team 'Innovators of the Year 2025'.

The technology emerging from LTU reflects deep European research excellence rooted in adaptive hardware systems, runtime orchestration, advanced systems engineering, embedded systems, and energy-efficient computing design. D-Chiplet combines these strong Swedish academic foundations with entrepreneurial and commercial experience shaped in Silicon Valley. The company brings together world-class European research with leadership experience from the global semiconductor and compute ecosystem, creating a bridge between cutting-edge

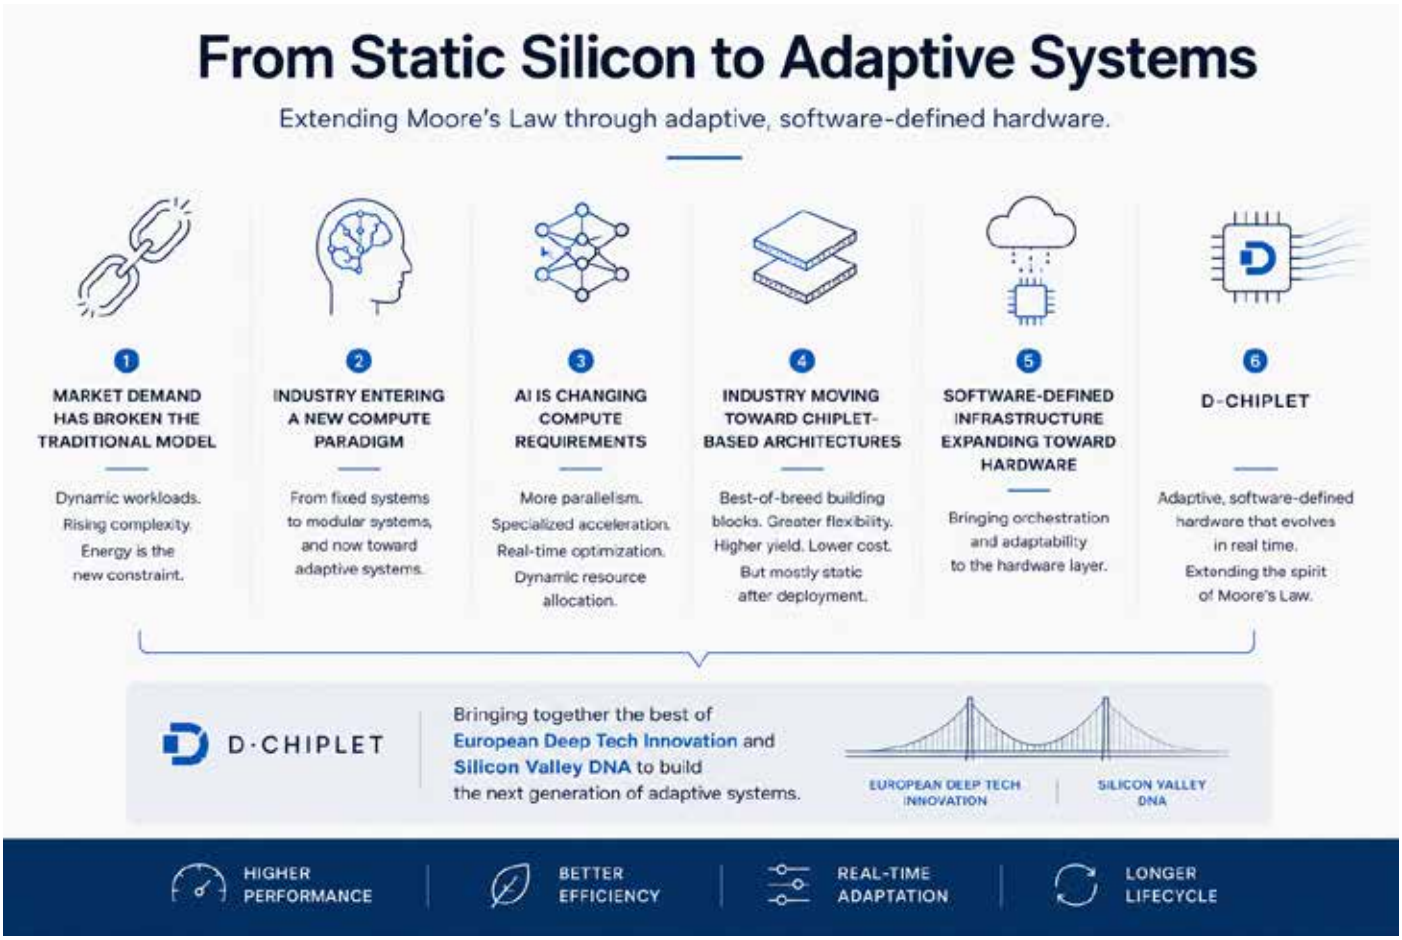

European innovation and the commercial scaling mindset often associated with Silicon Valley. Alongside the LTU research team, D-Chiplet also includes three experienced technology entrepreneurs, Daniel Hagström, Fredrik Berglund, and Keith Shea. Combined they have decades of experience building and scaling high-tech companies internationally. Daniel and Fredrik have both successfully launched more than a dozen venture-backed startups across the globe, and Keith brings more than 20 years of experience as a former executive of Intel Corporation.

In this sense, D-Chiplet represents more than a university spinout. It represents the combination of deep Swedish technical innovation with Silicon Valley entrepreneurial DNA. This research is now transitioning into industry through the formation of D-Chiplet AB, a company created to commercialize adaptive, software-defined hardware for heterogeneous chiplet-based systems.

D-Chiplet is building a runtime control and orchestration layer for heterogeneous chiplet-based systems. The simplest way to think about the platform is this: D-Chiplet brings a "microservices model" to silicon. Rather than relying on fixed hardware configurations, the platform enables compute resources to be dynamically orchestrated across heterogeneous hardware elements based on real-time workload requirements. At the core of the architecture is the concept of treating compute, memory, accelerators, and I/O resources as composable elements that can be dynamically assembled into functional "microsystems."

This enables dynamic allocation of compute resources, workload-aware optimization, and adaptive balancing of performance and energy consumption. It can also support longer hardware lifecycles as workloads evolve.

Unlike traditional instruction set architectures (ISAs), which define how software interacts with fixed hardware abstractions, adaptive systems allow the hardware resources to be reorganized in response to system requirements. Unlike virtualization layers, which partition existing resources, adaptive orchestration enables hardware resources to be composed and reconfigured for changing workloads. And unlike conventional operating systems, which manage software processes, D-Chiplet introduces a runtime control plane for orchestrating hardware.

The result is a transition from static hardware building blocks toward a fluid compute fabric.

Why this matters for Europe

Beyond its technical significance, this transition carries strategic importance for Europe. As global competition intensifies across semiconductors, AI infrastructure, and digital sovereignty, Europe has a significant opportunity to define new layers of the compute stack rather than competing solely on manufacturing scale. This creates an opportunity for Europe not only to participate in the next era of computing, but to help define critical architectural layers within it.

Europe already possesses deep strengths in:

- embedded systems,
- industrial automation,
- cyber-physical systems,
- telecommunications,
- automotive engineering,
- and energy-efficient computing.

Adaptive, software-defined hardware aligns directly with these strengths. The emergence

of heterogeneous chiplet ecosystems also creates new opportunities for European companies and research institutions to contribute differentiated intellectual property and system-level innovation. In this context, D-Chiplet represents more than a single technology platform. It reflects a broader transition in computing architecture, one where orchestration, adaptability, and energy efficiency become as important as raw transistor scaling. This is especially relevant as Europe seeks to strengthen its role in resilient digital infrastructure, AI competitiveness, and strategic autonomy.

Applications across strategic industries

The implications of adaptive hardware extend across multiple industries.

Automotive

Modern vehicles are evolving into software-defined platforms integrating autonomous driving, connectivity, infotainment, and safety-critical systems. Adaptive compute architectures can dynamically allocate resources based on changing operational requirements while improving energy efficiency and system longevity. For example, compute resources could be prioritized differently during highway autonomy, urban driving, charging, or diagnostic operation.

Edge AI and robotics

Edge systems operate under strict constraints on latency, thermal limits, and power consumption. Adaptive orchestration allows systems to optimize compute resources dynamically as environmental conditions change.

Telecommunications

As networks evolve toward distributed 5G and future 6G architectures, adaptive compute enables more efficient management of variable and distributed workloads. Examples include dynamic allocation of acceleration resources for baseband processing, edge inference, traffic management, or network slicing.

Data centers

Cloud infrastructure increasingly depends on heterogeneous acceleration and energy optimization. Reconfigurable hardware systems can improve utilization rates and performance-per-watt while reducing operational costs.

Aerospace and defense

Mission-critical systems require both determinism and adaptability. Dynamically orchestrated hardware can improve resiliency while supporting evolving mission requirements over extended lifecycles.

Toward the next era of computing

The transition from static silicon to adaptive systems represents more than an incremental improvement in semiconductor design. It signals the emergence of a new computing paradigm. For decades, the industry focused primarily on scaling transistor density and increasing raw performance. In the coming era, adaptability, orchestration, and energy efficiency will become equally important dimensions of compute architecture.

This transition cannot be solved by any single company or institution alone. It requires collaboration across the semiconductor ecosystem, including research institutions, system integrators, software developers, chip vendors, infrastructure providers, and industrial users. That collaborative ecosystem is precisely the kind of environment organizations such as INSIDE are helping to build.

The future of computing will not be defined solely by faster chips or smaller process nodes. It will increasingly be defined by systems capable of adapting intelligently to changing workloads, operating conditions, and energy constraints in real time. In that transition, adaptive, software-defined hardware may become one of the foundational architectural shifts of the AI era.

Keith Shea
Co-Founder D-Chiplet AB
Stockholm, Sweden

Keith Shea, based in Stockholm, Sweden, is an entrepreneurial technology executive with extensive international leadership experience building, scaling, and operating global businesses at the intersection of artificial intelligence, semiconductors, cloud software, and mission-critical systems. He has held senior executive roles across the United States and Europe, including more than two decades at Intel Corporation.

In addition to his operating roles, Keith serves on the boards of several commercial and non-profit organizations and is actively engaged in advising early-stage technology ventures. A native of Silicon Valley, he now resides in Europe with his family.

Keith holds a Bachelor's degree in Economics from Boston College and an MBA from The Wharton School at the University of Pennsylvania

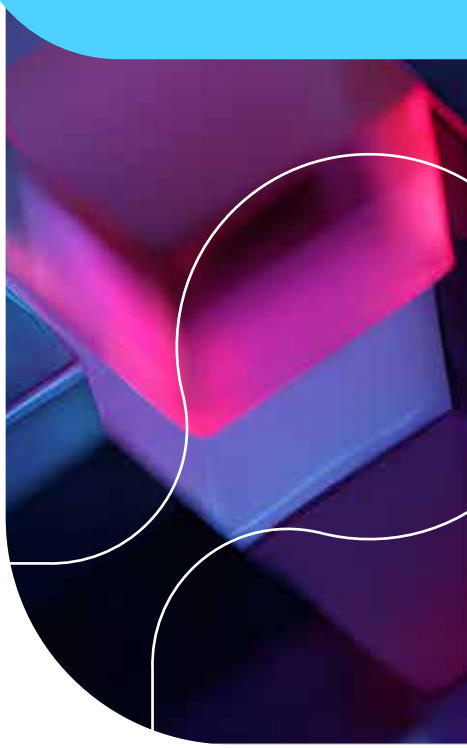

Chips JU

WECS 2026

2&3 December
Helsinki

SAVE THE DATE!

Aeneas

EPoSS. European Association on Smart Systems Integration

INSIDE Industry Association

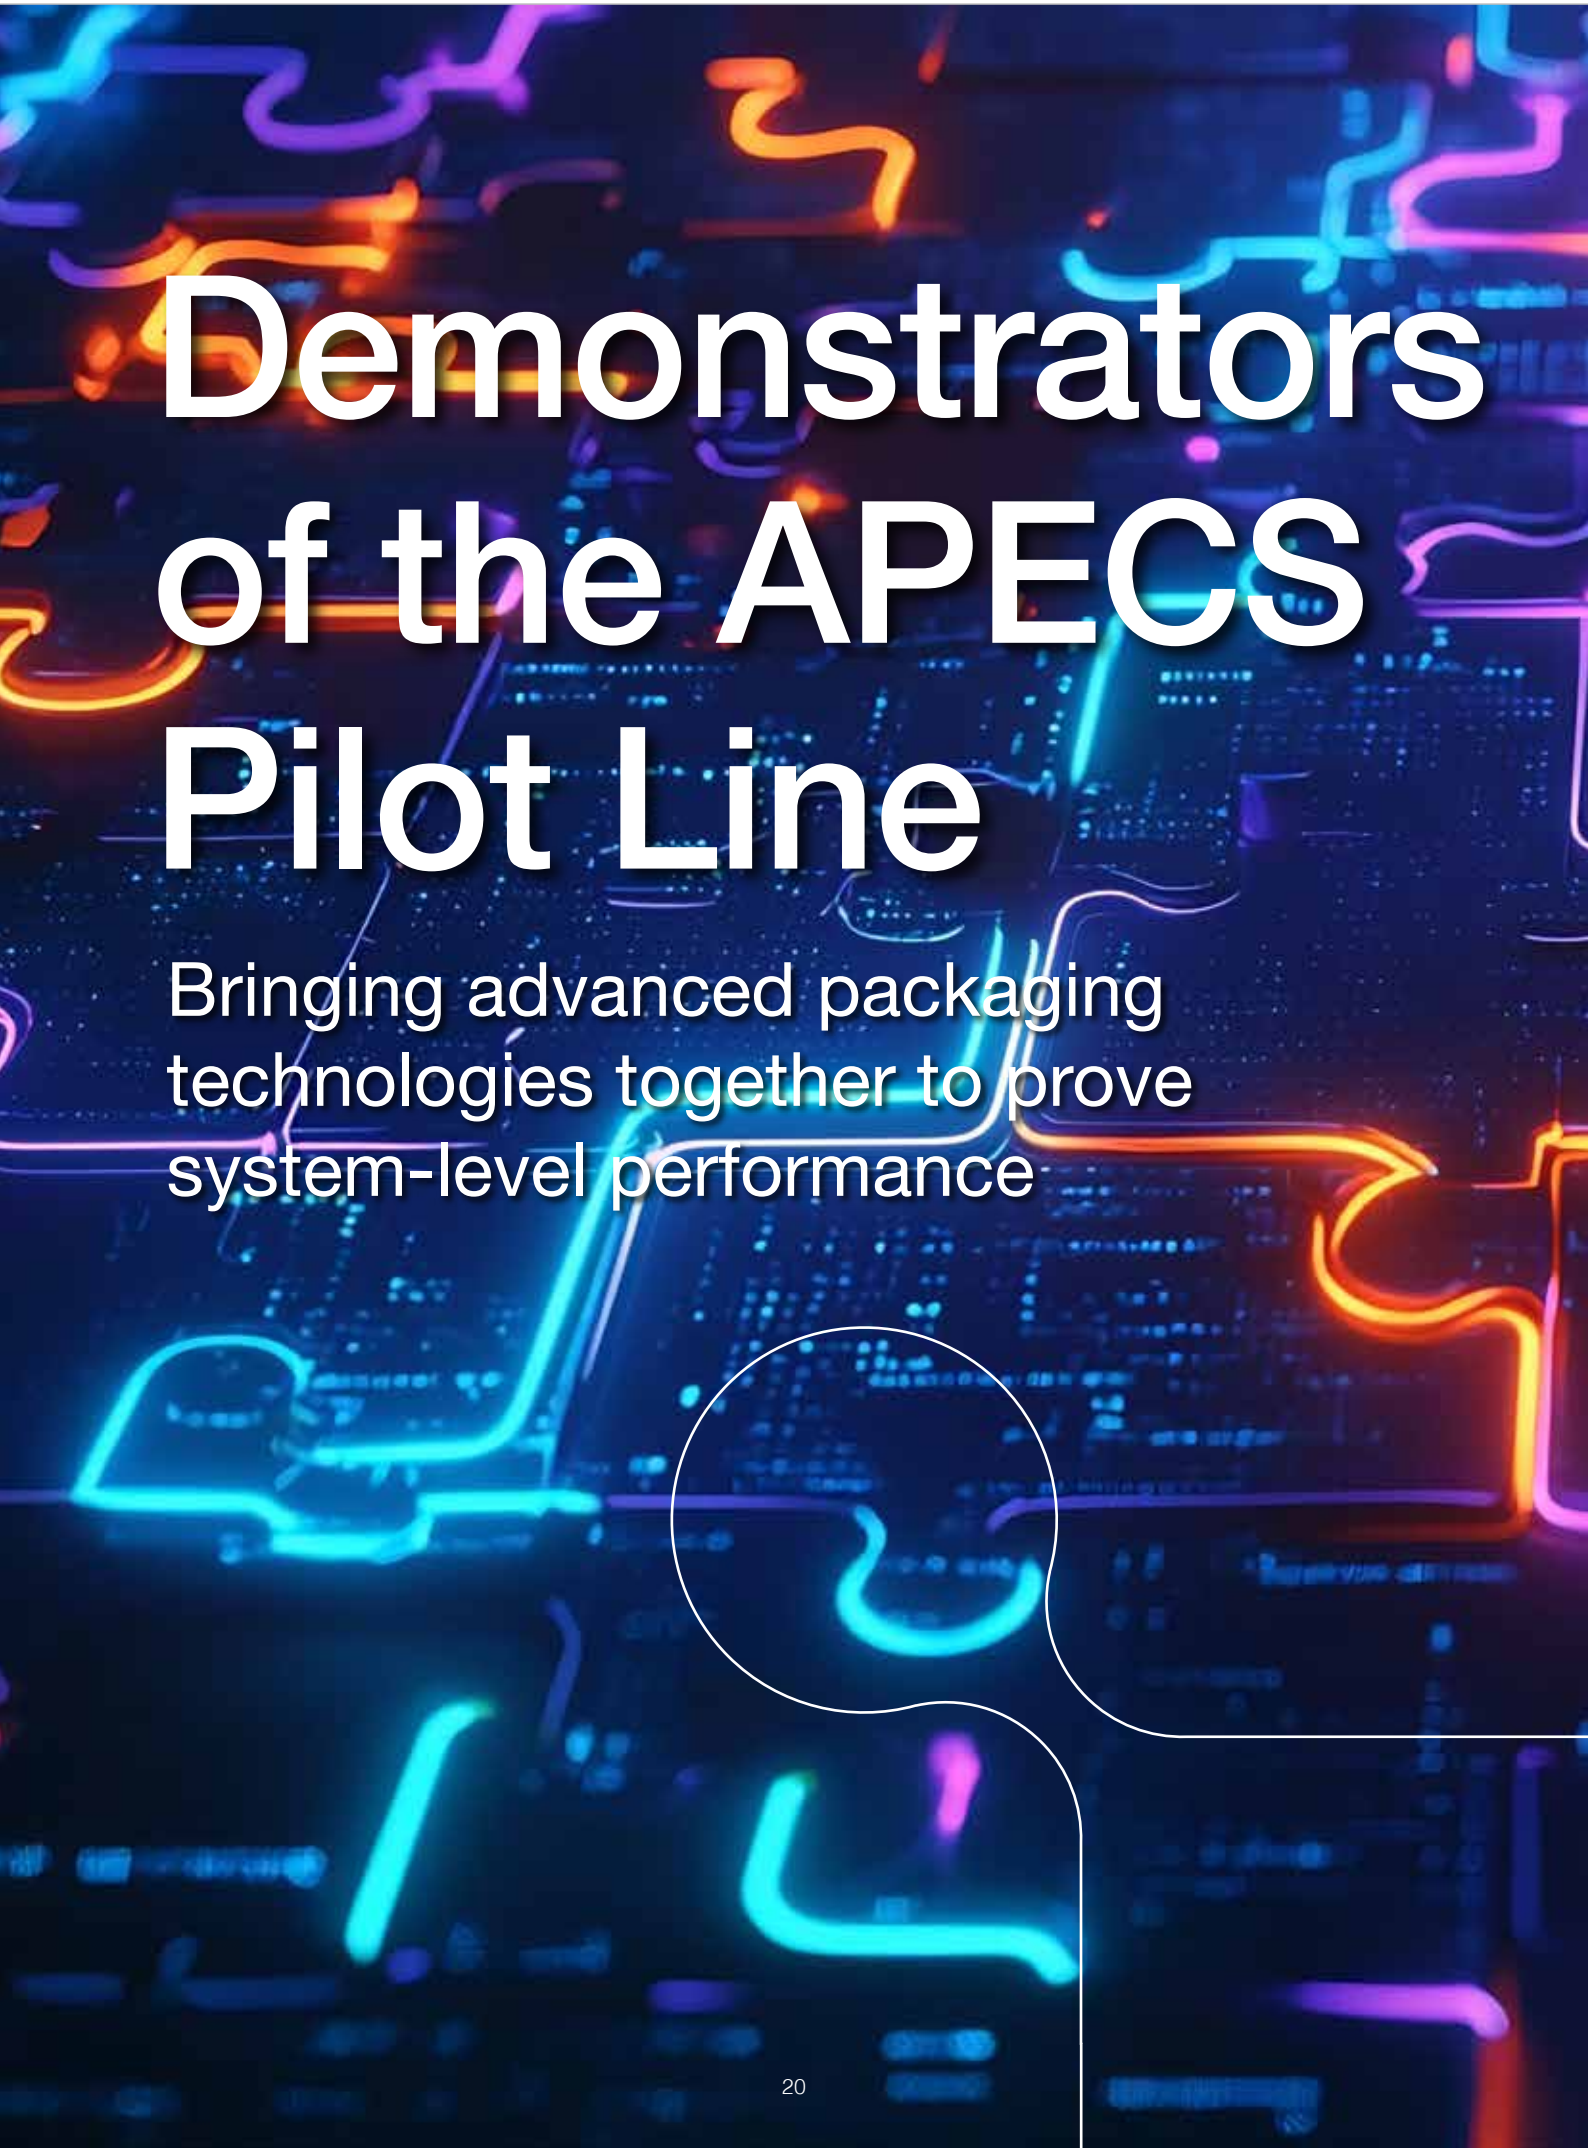

Demonstrators of the APECS Pilot Line

Bringing advanced packaging technologies together to prove system-level performance

Dr. Dirk Schumann
Director APECS Pilot Line
Research Fab
Microelectronics Germany
(FMD)

In advanced microelectronics, the competitive edge is no longer defined by individual technologies, but by the ability to integrate them into reliable, scalable systems. As chiplet-based architectures and heterogeneous integration become the new standard, the key challenge lies in translating fragmented capabilities into coherent production flows. The APECS Pilot Line addresses this shift by placing demonstrators at the center of its approach: not as showcases, but as integration testbeds that validate how technologies interact under real industrial conditions and ultimately deliver measurable system-level performance.

Europe's ambition to strengthen its microelectronics ecosystem increasingly hinges not only on mastering individual technologies, but on the ability to combine, integrate, and validate them at system level. As heterogeneous integration, chiplet-based architectures, and advanced packaging continue to reshape how electronic systems are designed and manufactured, the critical bottleneck is no longer access to single process steps, but the functionality of entire production chains across technologies, sites, and disciplines.

This is precisely where APECS – the European Pilot Line for Advanced Packaging and Heterogeneous Integration for Electronic Components and Systems – positions itself. Rather than focusing on a single technology or fabrication node, APECS has been conceived as a pan-European integration platform: bringing together distributed expertise, advanced infrastructure, and system-level design methodologies to enable next generation heterogeneous integration systems.

A pan-European pilot line built for integration

APECS is structured as a decentralized pilot line, pooling the technological capabilities of ten partners across eight European countries. Its foundation lies in the recognition that next-generation electronic systems, spanning high-performance computing, photonics, radio frequency, sensing, or safety-critical applications, cannot be realized within isolated process silos. Instead, they demand seamless interaction between materials, components, design flows, and production environments.

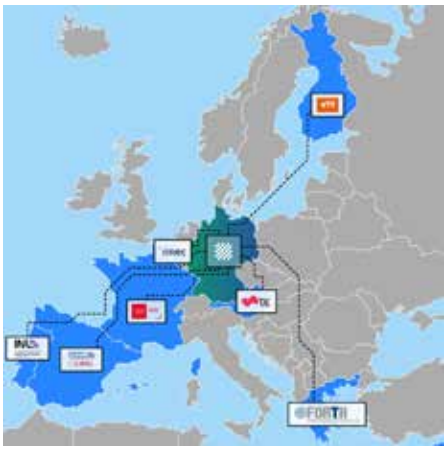

Figure 1: APECS Consortium: Germany (Fraunhofer-Gesellschaft as coordinator, FBH, IHP), France (CEA-Leti), Belgium (imec), Finland (VTT), Austria (TU Graz), Greece (FORTH), Spain (IMB-CNM, CSIC) and Portugal (INL).

What makes APECS distinct from other pilot lines is this explicit focus on integration rather than fabrication alone. Building on the infrastructure of the Research Fab Microelectronics Germany (FMD), APECS connects multiple specialized cleanrooms and process lines into a coherent, end-to-end production environment. A manufacturing execution system (MES) ensures consistency and quality across sites, while a central one-stop-shop-office (OSSO) coordinates access, requirements, and solution pathways for industrial as well as academic users.

In doing so, APECS reflects a shift in how technological value is created in microelectronics: away from isolated technological advancements and toward co-optimized systems, where design, packaging,

testing, and reliability are addressed collectively.

Technological scope: beyond classical advanced packaging

At the technological core of APECS lies a comprehensive platform for 2.5D and 3D heterogeneous integration, encompassing an unusually broad range of component technologies. These include CMOS, SiGe (BiCMOS) and as well as III–V-based devices for radio frequency, photonics, sensing, and emerging mixed-signal applications. Rather than treating these technologies as parallel options, APECS aims to combine them within shared architectures, enabling new system concepts that go beyond established packaging standards.

This ambition is reflected in four tightly linked technological focus areas:

- Quasi-Monolithic Integration (QMI): integrating multiple semiconductor functions into a single, compact solution without requiring all components to be fabricated on exactly the same silicon die or process node.
- Chiplet integration platforms: supporting both 2.5/3D integration technologies and allowing flexible assembly of modular, application-specific systems.
- Characterization, testing, and reliability (CTR): a holistic approach that combines characterization and testing as well as key innovation for safety and functional reliability, enabling continuous validation across the production chain.
- System-Technology Co-Optimization (STCO): a design framework that explicitly connects system requirements with technology choices, process flows, and manufacturing constraints, resulting in faster, more cost-effective, high-performance solutions for next-generation semiconductor products.

Together, these focus areas address a central challenge in advanced packaging: Ensuring that heterogeneous systems are not only technically feasible, but manufacturable, reliable, secure, and scalable.

From technology modules to system performance

A defining feature of APECS is that its technological focus areas do not exist in isolation. Development in QMI, chiplet platforms, CTR, and STCO are continuously aligned with system-level integration goals. The intent is not to optimize individual process steps, but to understand how decisions at one stage, materials, interconnects, layouts, or assembly, affect performance, yield, and reliability across the entire production chain. This approach is particularly relevant for chiplet-based systems, where interfaces exist not only at material and electrical level, but also in design data, testing strategies, and logistics. By explicitly addressing these interfaces, APECS creates a framework in which gaps, risks, and trade-offs become visible early, long before technologies are transferred to manufacturing.

In practical terms, this means that technological progress within APECS is measured not only by process maturity, but by how well different elements can be combined into functioning, repeatable system flows.

The crucial role of demonstrators within APECS

This systemic perspective explains why demonstrators play a central role within APECS. They are not conceived as showcase end products, nor as isolated proofs of concept. Instead, demonstrators serve as integration checkpoints – vehicles through which the full pilot line is exercised, tested, and refined. To this end, four demonstrators are being developed within APECS to identify interface incompatibility, process gaps, and data-flow issues that would remain otherwise hidden.

These demonstrators are designed to validate the achievements of the technological developments across multiple application domains. Each of the four demonstrators addresses a key field of the APECS pilot line. In doing so, they make use of the pilot line's decentralized structure and its STCO framework to assess how technologies can be jointly deployed across distributed infrastructures.

A particular focus lies on the logistical and technological interfaces between different cleanroom environments, each with its own requirements, as well as on enabling designers to access and combine diverse technologies. The objective is not only to optimize the transfer between individual process steps, but also to enable new process combinations that extend the capabilities of chiplet technologies and heterogeneous integration, while maintaining efficiency levels comparable to a centralized production environment.

Equally important, demonstrators establish a shared reference point for interaction with industry. They ground discussions with industrial stakeholders in observable system performance rather than abstract capabilities, while creating a feedback loop to identify and transfer requirements of new applications to the technological roadmap. For industry and investors alike, demonstrators provide tangible evidence of performance, reliability, and application relevance making the integrated capabilities of the APECS pilot line both visible and accessible.

In this way, APECS functions as a bridge between technological capability and system-level integration. Its core strengths – ranging from quasi-monolithic and 2.5/3D integration, chiplet design platforms and advanced characterization, which all is enclosed by STCO – provide the new technological backbone for Europe.

Demonstrators as integration frameworks

Four demonstrators are developed within the APECS pilot line, each addressing different key challenges in microelectronics:

1. The High-Performance Computing (HPC) demonstrator evaluates System-Technology Co-Optimization (STCO), where customer requirements such as cost, size, and energy consumption are built directly into chiplet and chip design e.g. for data center or edge AI solutions.

2. The Multi Materials Sensor (MMS) module showcases the modularity of chiplets by integrating diverse sensors such as gas sensors, optical, and acoustic systems, while managing different thermal expansion, heat dissipation and cross interference.

3. The Photonic Integration (PI) demonstrator focuses on edge InP technologies and photonic wire bonding to increase bandwidth and reduce energy consumption, addressing the growing demand for high-speed data transfer in compact systems.

4. The Radio Frequency (RF) demonstrator will show innovative solutions for both wireless data communication and radar applications. The focus is on enhancing performance by combining the advantages of various semiconductor technologies, which are integrated using advanced heterogeneous integration techniques.

Rather than functioning as prototypes of end products, they serve as platforms to validate chiplet-based process chains across the pilot line. Together, these demonstrators show how APECS addresses the complexity of future microelectronic systems by developing optimized, scalable, and reliable solutions.

High performance computing (HPC) demonstrator

The accessibility to design and manufacture customer-tailored systems for High Performance Computing (HPC) and edge AI applications is extremely limited in Europe. While most of the required services are individually available within Germany, they widely lack an incorporation into a holistic solution's landscape. The pilot-line's HPC demonstrator implementation will take a subset, leveraging System-Technology-Co-Optimization (STCO) to build two chiplet-based modules for Data Center and Edge AI systems.

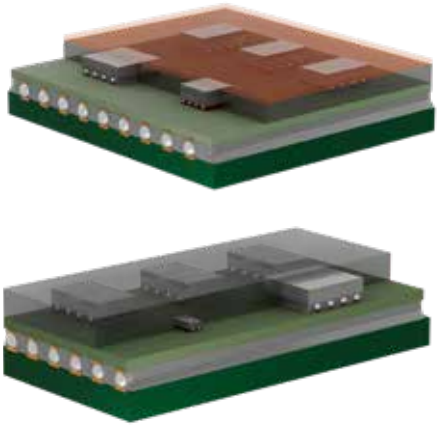

Figure 2: Conceptual illustration of chiplet-based modules for Data Center with heat sink on top (left) and Edge AI systems without heat sink (right).

STCO is about integrating already available services such as

- system design for high-density integration, considering available standards like UCIe –S and BoW for chiplet-to-chiplet interconnections and new solutions
- System-on-chip (SoC), chip and chiplet design
- modelling and simulation of the physical module implementation, assembly, and physical integration of the modules
- module/system test and validation

These services are meshed together and completed by others to enable a seamless and concurrent interoperability, to fill integration gaps and overcome interface challenges within the full production chain, both at the material and data levels.
For customer projects the pilot line will also offer reliability & product qualification. Eight Fraunhofer-institutes are involved in this demonstrator in chiplet architecture specification, fabrication as well as test, packaging development, interposer and assembly developments: IIS & AISEC, IZM, IZM-ASSID, IPMS & FHHG IIS, IIS-EAS & AISEC.

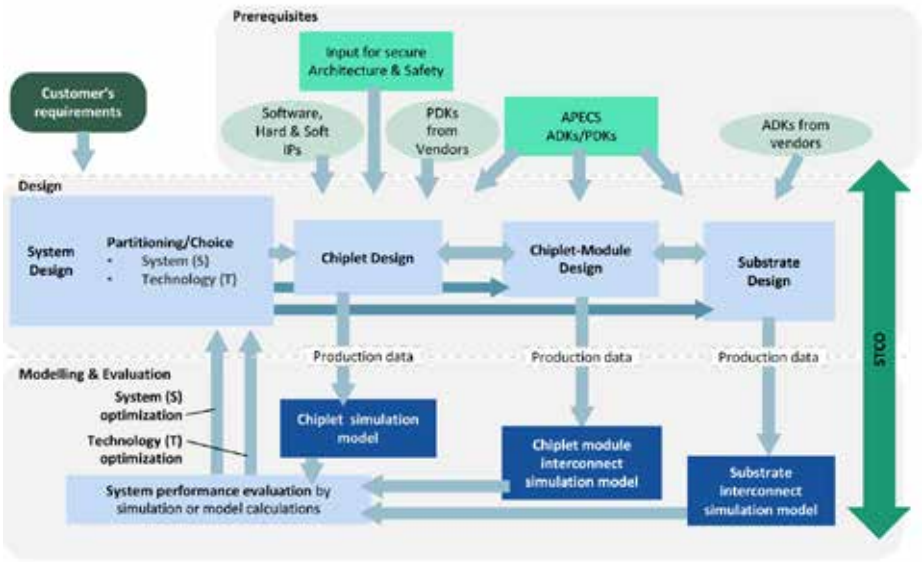

Figure 3: Sketch of the part of the holistic APECS System-Technology Co-Optimization (STCO) process that applies to the design of the HPC demonstrator modules.

The STCO process is required to provide valid solutions with reasonable effort. The main reasons are:

- The partitioning and technology choice decisions interfere with each other in terms of system performance and budgetary implications as well as manufacturability, market and supply chain considerations, since optimizing only one after the other will not find an adequate solution.
- Changes in technology decision result in changes of electrical properties and the outcomes in terms of system performance are not always obvious - therefore, several options need to be evaluated by modelling and simulation of the parts.
- While the inference between multiple parts is reciprocal, the parts need to evolve synchronously together with a simulation model describing the interaction for the opposite part.
- There is only a guess what the performance will be after the design, but the real performance will be available only after extracting the design data into models and calculating performance from the extracted models.

To address these and more requirements, the said STCO flow was established. Customer requirements and design prerequisites feed the concurrent design of the parts. Then the system performance is calculated as in the current iteration's design data described (Figure 3).

When not all the customers' requirements are met at first, a new iteration of the Co-Optimization loop needs to be started with an idea of how to change the system or to negotiate the requirements in case the fulfilment is too costly or impossible.

This procedure is validated by planning two different systems: One with compute power for industrial computing and data centres and one with compute efficiency for Edge-AI. Both systems make use of organic package substrates, which basically enable bandwidth densities around 100 GB/s/mm^2 according to the UCIe-S standard.

UHD silicon interposer push the edge of feasibility

Addressing the future needs for significantly greater bandwidths, the involved pilot line's partners strive to adapt the STCO flow for ultra-high density (UHD) heterogeneous integration on 300 mm Si wafer level. Since the interfaces between foundry-manufactured wafers and the integration of respective chiplets into next gen HPC systems lack the required rigorous definitions, the targeted technology demonstrator has no compute

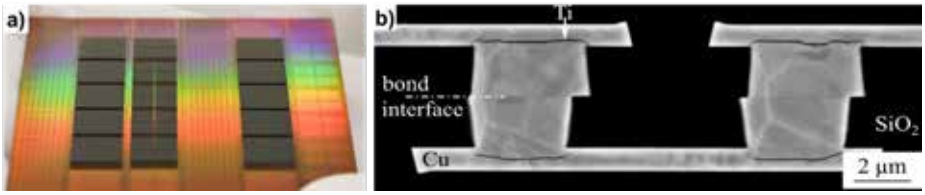

Figure 4: Section (Coupon) of a 300 mm wafer with hybrid bonded chips (a) and cross-section of the hybrid bonding interface (b).

function but will implement several leading-edge advanced packaging technologies, such as

- Ultra-high density redistribution layers with integrated capacitors and resistors to support signal integrity and power supply.
- High-density mixed-pitch interconnections based on hybrid bonding (Figure 4) and microbumps to enable bandwidths above 1 TB/s/mm² according to UCIe-A and UCIe-3D.

Multi-Material Sensors (MMS)

One platform, many sensors: The case for a universal silicon interposer

The MMS demonstrator of the APECS pilot line integrates GaN gas sensing, acoustic CMUT, opto-chiplets and high-resolution magnetic Hall-arrays on a reusable interposer – a modular approach that lowers barriers for sensor development across industry.

The problem worth solving

Sensors rarely exist in isolation. In practice, whether for industrial monitoring, medical diagnostics or environmental analysis, useful sensing systems need to detect more than one quantity, often using different physical or chemical principles. The technologies behind these sensors, gallium nitride, silicon MEMS, photonic waveguides, capacitive transducers, magnetic Hall structures, come from very different material and process environments. Today, integrating them into one system almost always means starting the packaging design from scratch.

The MMS demonstrator within the APECS pilot line is a direct response to that problem. Its goal is not to develop one new sensor, but to show that a common silicon interposer infrastructure can serve as the integration backbone for several sensor technologies simultaneously, and that this approach can work at an industrially relevant scale.

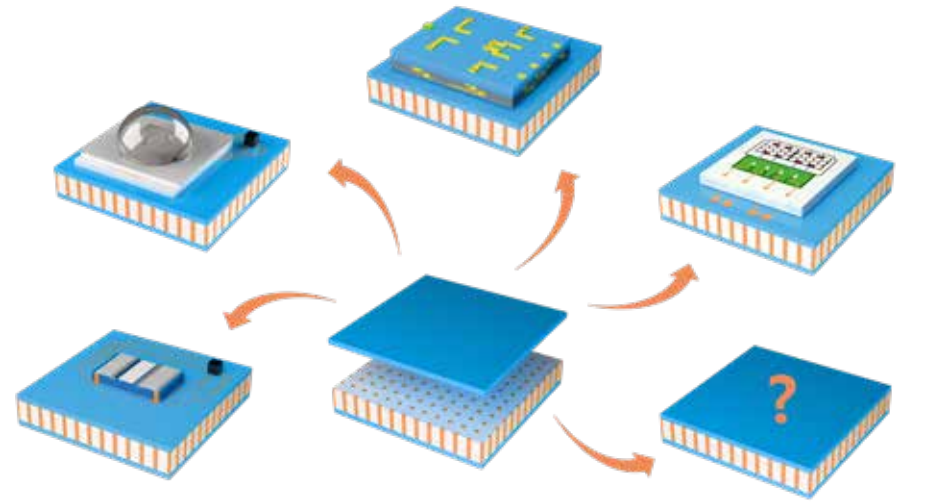

Figure 5: A universal silicon interposer connects multiple heterogeneous sensor chiplets. The open slot (question mark) represents a position available for future or customer-specific sensors.

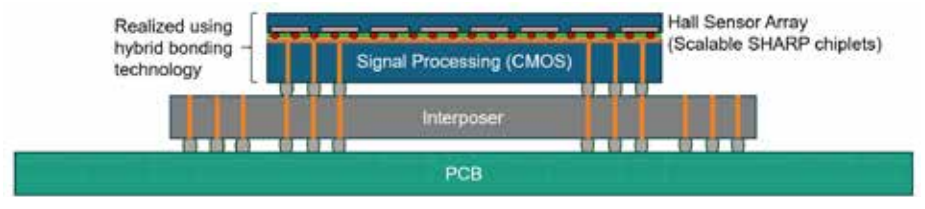

Figure 6: SHARP magnetic sensor chiplet: monolithic Hall-array on silicon hybrid-bonded to CMOS, tiled on a TSV interposer for scalable high-resolution magnetic field measurement illustrating the universal interposer stack.

A platform, not a package

At the centre of the MMS concept is a silicon-based interposer with a universal core design. This is the component that all sensor chiplets are built around. The interposer features very deep through-silicon vias in the range of 300 to 400 µm. What makes the platform reusable is the redistribution layer (RDL): rather than designing a new interposer for each sensor combination, only the RDL is adapted to the specific interface geometry and signal requirements of the respective chiplet.

This separation of stable backbone and adaptable interface is what the APECS team refers to as a System-Technology Co-Optimization (STCO) approach. All sensor chiplets must comply with a defined interface standard. The analogy is closer to a standardized connector system than to a conventional custom package, and it has direct consequences for how quickly new technologies can enter the platform and reach industrial use. The open slot with the question mark in the interposer figure is therefore more than a graphic element, it is intended to indicate that additional sensors can be added to the same backbone.

Four sensor families, one platform

The MMS demonstrator currently integrates four distinct sensor technologies, each representing a different application domain. The first is a GaN-based HEMT hydrogen gas sensor designed for operation in harsh environments. A GaN-on-Si/QST sensor chip is heterogeneously integrated with a silicon readout IC on the interposer. GaN HEMT structures are well suited to this application because of their thermal stability and sensitivity to surface charge changes induced by gas adsorption. First chiplets from this sub-demonstrator have already been fabricated, and the development target is the first 8-inch European GaN sensor platform, which is intended to be accessible to SMEs.

The second technology is an acoustic sensor (Si-CMUT). Several CMUTs are co-packaged with a CMOS chip on a silicon pocket wafer using QMI technology, and the resulting chiplet is then mounted onto the common interposer. Preliminary process studies for this

chiplet have been completed successfully. The chiplet makes the acoustic integration accessible to SMEs in the sensor market.

The third set of components is the opto-chiplet, which consist of two separate devices. The first is a MEMS mirror based on an epi-poly-silicon device assembled in a hermetically sealed package with a glass interposer and a glass dome for optical access. The second is an electro-optical chiplet that uses post-CMOS photonics: QMI technology allows photonic waveguides and circuits to be fabricated on top of a CMOS wafer after the transistor process is complete. Both chiplets are mounted onto the interposer using the same platform logic as the GaN and acoustic components.

Another sub-demonstrator is the SHARP concept (Silicon Hall-Array High-Resolution Platform), which is based on a monolithic three-dimensional high-resolution Hall sensor array in silicon that is stacked onto a separate CMOS substrate for signal processing using hybrid wafer-to-wafer bonding. Each SHARP chiplet behaves as a self-contained measurement system, but several chiplets can be placed side by side on the interposer to form a seamless, large-area magnetic field sensor array. Typical application scenarios include magnetic field cameras, non-destructive testing in mechanical engineering, and field homogeneity measurements.

Together, these four sensor families show that the MMS interposer is not tied to a single device type.

Why this matters and what it changes

The competitive argument for the MMS platform rests on a straightforward observation: packaging design today is largely technology-specific. A company that develops a GaN-based device faces a different packaging workflow than one working with acoustic MEMS and combining both in one product typically means building a new assembly architecture from the ground up. Adding the opto-chiplets or high-resolution magnetic arrays make the situation even more complex. This is slow, expensive and creates a structural disadvantage for smaller players who lack the resources to repeatedly qualify new packaging processes.

The MMS interposer concept addresses this at the architecture level. By defining a common interface standard and a reusable core design, it allows new sensor chiplets to enter the platform without triggering a full system redesign. The question-mark slot in the interposer figure is the practical expression of this: it is a real integration position, not a placeholder, and it is meant to communicate to industrial partners that the platform is open for extension.

Finally, the APECS approach is explicit about supply-chain accessibility. The demonstrator work includes interface standardization and supply-chain coordination across all sub-demonstrators, with the specific aim of making the platform usable by SMEs. The first 8-inch GaN sensor platform being developed within the HEMT sub-demonstrator carries this intent directly: it is designed from the start to be available to companies that do not have in-house access to GaN wafer-scale processing. The chiplet format of the CMUT, opto- and SHARP modules follows the same logic, advanced sensor functions are provided as standardised building blocks that can be

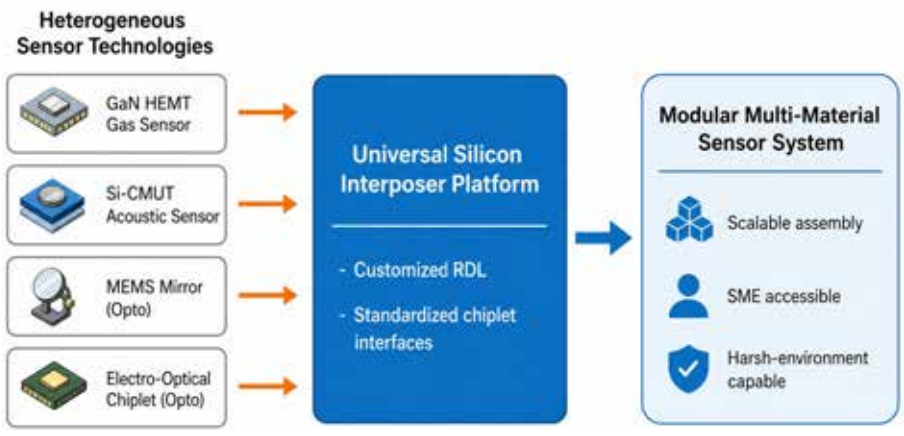

Figure 7: Platform integration logic: heterogeneous sensor technologies share a common interposer backbone, reducing redesign effort and accelerating industrial transfer.

combined on the interposer without each user having to develop their own process chain.

Taken together, the MMS demonstrator of the APECS pilot line shows that heterogeneous sensor integration does not require a different approach for every material class. A well-designed interposer platform with standardized interfaces and an adaptable redistribution layer can serve as the shared infrastructure for GaN, acoustic, opto-chiplets and magnetic Hall-array technologies simultaneously. The result is a modular sensor system platform that reduces integration complexity, shortens time to market and opens advanced packaging to a broader industrial user base.

Photonic integration (PI)

With the Photonic Integration demonstrator, APECS pilot line's capability regarding hybrid-integration of III-V opto chiplets with III-V electronic chiplets will be demonstrated:

A high speed 1300nm InP based Electro-absorption Modulated Laser (EML) 4 channel array will be hybridly integrated with an InP based DHBT driver chiplet on a common based interposer including all electrical high-speed connections and coupling to an optical fiber array. Target is to achieve a 1300nm 4 x 200Gbps PAM4 transmitter (schematic view see Figure 8)

A comprehensive transmitter hybrid integration concept was developed, based on a Si-based interposer that hosts the EML 4-array chiplets, the InP driver chiplets, and the associated highfrequency RF feeding lines. This interposer will be cointegrated with a fiberarray unit on a suitable heat dissipative submount. The required metallization and insulation layers have been defined and currently work focusses on the detailed RF line layout and the mounting and alignment processes and sequence for the implementation of EML, driver chiplet and fiber array.

A major achievement of the first project year is the reduction of the InP EML 4-array pitch from 640 µm to 375 µm, significantly increasing the shoreline density while simultaneously improving device performance: the EML modulation speed could be increased from 200 Gbps PAM4 to 290 Gbps PAM4. The achieved shoreline density currently is a record value for such EML 4-arrays. The final target in APECS is to achieve a pitch of 250 µm only.

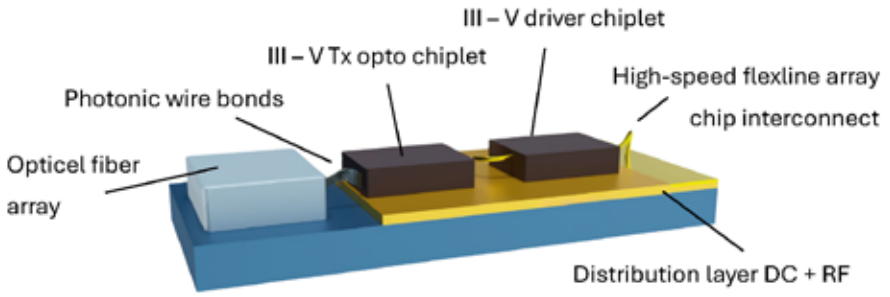

Figure 8: Schematic cross section of the tackled high speed optical transmitter.

Progress was also achieved on the InP DHBT driver chip, which now reaches an electrical bandwidth of more than 160 GHz. Further speed improvements are planned.

The optical interface between the EML 4-array and the fiber 4-array was further advanced, with the first successful demonstration of multichannel optical coupling using photonic wire bonds, a key enabler for alignmenttolerant coupling and for bridging disparate modefield diameters.
The PI-Demonstrator transmitter represents an example for the hybrid integration of III-V optoelectronic components on a common interposer inclusive fiber coupling. The developed technology will pave the way for the future realization of customized hybridly integrated photonic and electronic chiplets within the APECS pilot line.

Radio frequency integration (RF)

The RF integration demonstrators will showcase the potential of the APECS pilot line in the field of complex, highly integrated RF systems up to 325 GHz. The performance and innovative potential of various combinations of semiconductors and heterogeneous integration technologies will be demonstrated. The applications addressed range from next-generation high data-rate mobile communications and wireless links to advanced radar and sensing systems. Besides the various possible combination of advanced semiconductors by innovative heterogenous integration technologies, four RF technology demonstrators, as shown in Figure 9, will be realized: a InP-on-BiCMOS TRX: 6G D-Band Transceiver (1.), a Sub-THz BiCMOS-mHEMT Transceiver on Interposer (2.), a Flex D-Band Radar (3.) and a D-Band Communication Module (4.).

1. **InP-on-BiCMOS TRX: 6G D-Band Transceiver**
   The D-band (110…170 GHz) offers large absolute bandwidth for wireless communications as needed, e.g., for high-data rate point-to-point links in the backhaul network of 5G and 6G base stations. The heterogeneous integration of high performance SiGe BiCMOS chiplets together with bipolar InP chiplets will boost output power and efficiency

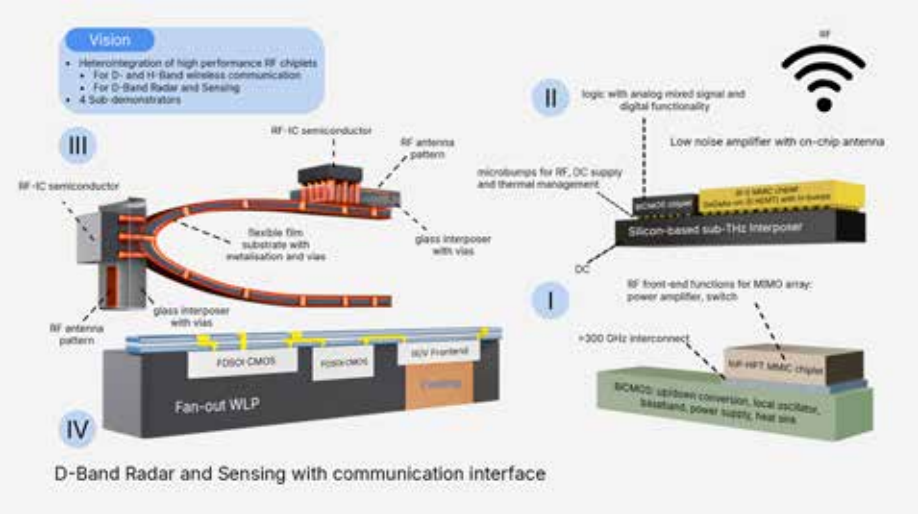

Figure 9: RF Demonstrator with four Sub-Demonstrators for wireless communication and radar applications.

at D-band, thus extending range and energy efficiency. The different building blocks and subcircuits of a transceiver will be partitioned between the SiGe-BiCMOS and the InP-HBT technologies considering the best performance of each individual technology to maximize the overall system performance. The approach provides the seamless integration of the two technologies with high performance and broadband interconnects using microbumps.

2. **Sub-THz BiCMOS-mHEMT Transceiver on Interposer**
   The Sub-THz BiCMOS-mHEMT Transceiver on Interposer demonstrator targets the heterogeneous integration of sub-THz chiplets and interposer technologies for high-resolution sensing above 200 GHz. The demonstrator aims to realize a high resolution FMCW radar frontend operating at 256 GHz by combining a SiGe BiCMOS transceiver chiplet together with InGaAs mHEMT amplifier chiplets on a common interposer platform. Glass and silicon interposers together with various flip chip interconnection technologies are applied to enable high-performance interconnects at such high frequencies. In summary, a scalable, low loss, compact and modular platform for sub-THz systems-on-interposer for sensing applications will be demonstrated.

3. **Flex D-band radar**
   Compared to sub-demonstrators 1. and 2., the Flex D-band radar will be a remarkable testcase of the pilot line's capabilities in integrating a multi-module radar system onto bendable and/or non-planar surfaces. The concept is based on a heterogeneously integrated D-band (110 – 170 GHz) radar transceiver mounted on an RF-glass interposer using flip-chip technology. In addition, broadband TX and RX antennas will be implemented directly on the RF-glass interposer and connected to the radar chiplet via low-loss interconnects. This approach offers significant advantages over on-chip antennas, including higher efficiency and wider frequency bandwidth. These radar modules – each consisting of a D-band frontend and an RF-glass interposer with an antenna array – will be integrated multiple times onto a flexible polymer interposer. The resulting innovation is a multi-module radar system that is well suited for integration into or onto bendable and/or non-planar surfaces, such as the wings of wind turbines or drones or moving robot arms to give just a few examples.

4. **D-band communication module**
   In alignment with sub-demonstrator 3. above, the sub-demonstrator 4. is also guided by the goal to show the integration of multiple chiplets. Here, a fan-out wafer level packaging (FOWLP) approach will be presented. FOWLP is well suited for RF and mmWave application due to shortest interconnects and feasibility of passive components and structure integration. The package will co-integrate two technologies: gallium nitride (GaN) and advanced FDSOI CMOS. Also cooling structures are integrated to ensure long-term stability of the integrated chiplets.

For each RF sub-demonstrator, the STCO (System-Technology Co-Optimization) process is used to achieve the highest possible system performance within a short development timeframe for highly integrated RF systems that combine different integration and semiconductor technologies. Within the APECS framework, this requires very close collaboration between Design and Design Enablement, as well as with semiconductor and integrations Technologies and the characterization, test and reliability capabilities and experts. In this sense, the demonstration projects are also the first "customers" to use the STCO and to conduct a proof-of-concept for the STCO.

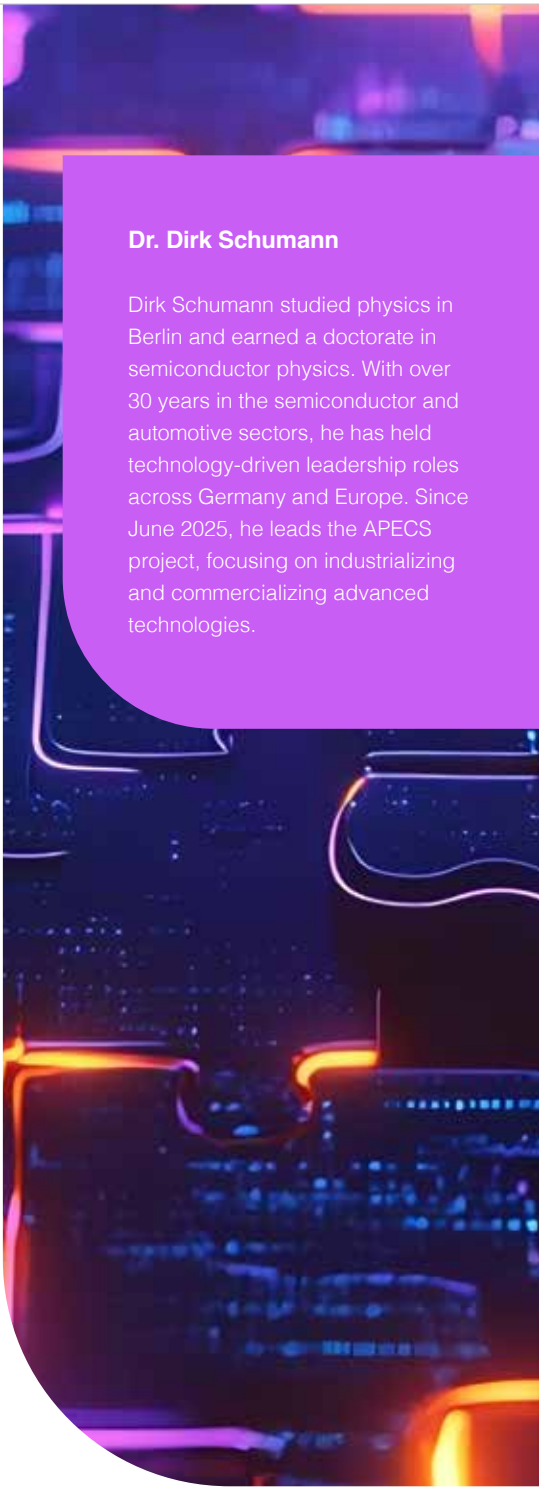

Dr. Dirk Schumann

Dirk Schumann studied physics in Berlin and earned a doctorate in semiconductor physics. With over 30 years in the semiconductor and automotive sectors, he has held technology-driven leadership roles across Germany and Europe. Since June 2025, he leads the APECS project, focusing on industrializing and commercializing advanced technologies.

APECS

Co-funded by

Chips JU

Co-funded by the European Union

With funding from the: Federal Ministry of Research, Technology and Space

Freistaat Sachsen · Baden-Württemberg · Berlin · Bayerische Staatsregierung · Land Brandenburg · SH Schleswig-Holstein · Sachsen-Anhalt

APECS is co-funded by the Chips Joint Undertaking and national authorities of eight member states within the framework of the EU Chips Act created under the „Chips for Europe" initiative of the European Commission. Overall funding for APECS amounts to € 730 million over 4.5 years.

Business Finland · CCDR Norte · anr · Flanders State of the Art · Hellenic Republic Ministry of Development · Federal Ministry Republic of Austria Climate Action, Environment, Energy, Mobility, Innovation and Technology · FFG

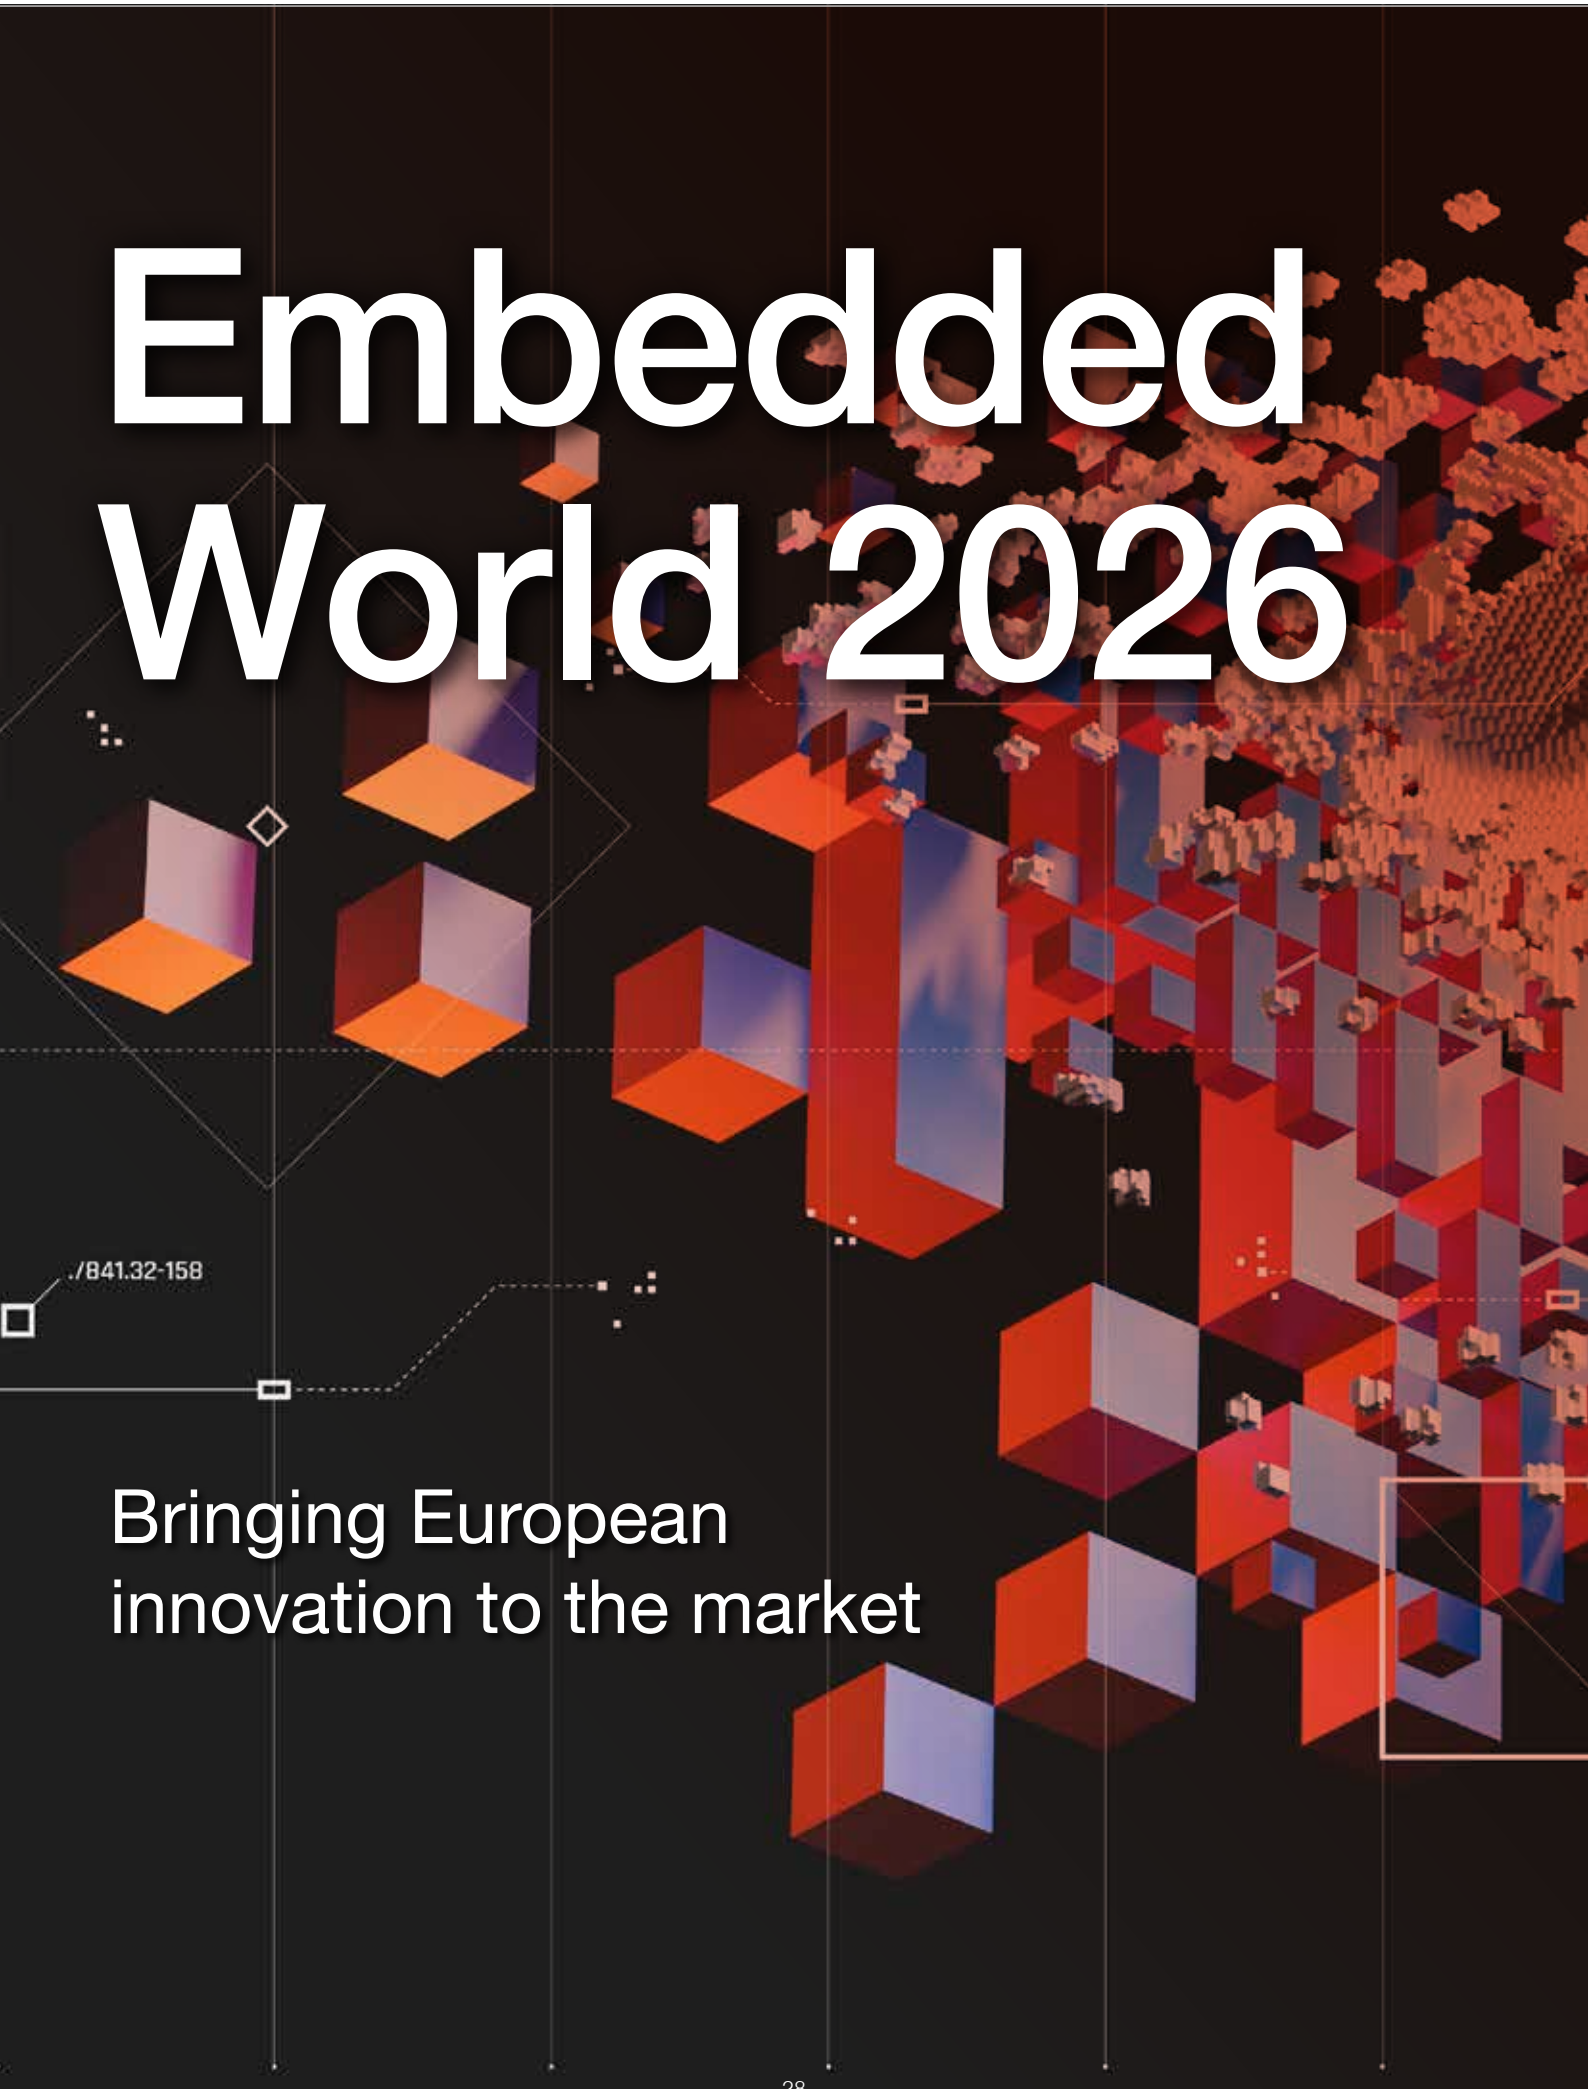

Embedded World 2026

Bringing European innovation to the market

Paolo Azzoni
INSIDE Industry Association

Maha Karim-Hosselet
INSIDE Industry Association

Santus Kisoka
INSIDE Industry Association

From 10 to 12 March 2026, the embedded systems market players gathered once again in Nuremberg for Embedded World, one of the world's leading trade fairs for embedded technologies. With more than 36,000 visitors, over 1,200 exhibitors, representatives from more than 90 countries and five exhibition halls covering the entire spectrum of embedded technologies, the event confirmed its position as a key global meeting point for industry, research and innovation

A global marketplace for embedded technologies

Embedded World is often described as an exhibition. In reality, it is much more than that. It is a marketplace where technologies compete for attention, credibility and adoption. Companies come to showcase products, identify customers, establish partnerships and assess their competitive positioning. Research organisations come to demonstrate that their innovations can address real industrial challenges.

This distinction is important. Unlike many research conferences, where projects often present results primarily to other researchers, Embedded World exposes innovations to a much broader audience of industrial stakeholders, technology providers, investors and potential customers. Bringing research and innovation results into such an environment requires confidence, maturity and a willingness to engage directly with market realities.

For Europe, this transition from research visibility to market visibility is essential. Technological excellence alone is not enough. Innovation creates value only when it reaches adoption.

Building a European presence

Embedded World 2026 marked the third consecutive participation of INSIDE Industry Association, its members and the Chips JU demonstrating a sustained commitment to strengthening Europe's visibility within the Electronic Components and Systems (ECS) ecosystem.

At the heart of this presence was a 128 m^2 pavilion bringing together sixteen European organisations from INSIDE community. The pavilion showcased expertise spanning chip design, artificial intelligence, embedded and cyber-physical systems, engineering automation, digital mobility solutions and other strategic technology domains.

More importantly, the pavilion represented far more than a collection of individual exhibitors. It provided a collective European showcase, demonstrating the richness, diversity and collaborative strength of the ECS community.

In an increasingly competitive global environment, visibility matters. Yet visibility alone is not enough. The real value of INSIDE pavilion lies in its ability to present Europe not as a set of isolated organisations, but as a connected ecosystem capable of delivering innovation across the entire value chain.

From visibility to impact

Throughout the three-day event, INSIDE pavilion became a busy spot for many demonstrations, networking activities and business discussions.

INSIDE members interacted with industrial stakeholders, research organisations, and potential customers. There were many targeted meetings that allowed discussing future cooperation, business developments and market opportunities. The live demonstrations allowed visitors to see first-hand the latest technologies developed by INSIDE members and their potential use.

These results were much more than just raising visibility.

INSIDE members identified many promising business opportunities, many new customer contacts and discussions about future cooperation. In addition, the event was

helping achieve some broader European goals related to technological excellence, industrial competitiveness and faster implementation of innovations from research environments into practical application.

What is probably the key message of Embedded World is that innovation acquires its value through feedback from the market. It is important for technologies working well in the research environment to show their relevance in industrial domains. Events such as Embedded World provide us precisely this opportunity.

This is an important step in transforming research investments into industrial impact.

Supporting European SMEs on the global stage

A particularly important objective of this initiative is to support European small and medium-sized enterprises (SMEs), which represent the backbone of the Electronic Components and Systems community. For many SMEs, participating independently in a major international trade fair such as Embedded World can be challenging. Beyond the direct costs, companies must compete for visibility in an environment dominated by large multinational corporations with substantial marketing resources and established global brands.

INSIDE Industry Association helps to address this challenge by providing SMEs with a highly visible and credible platform from which to showcase their technologies, products and services. By participating as part of a coordinated European presence, INSIDE members benefit from increased visibility, access to a broader network of stakeholders and the institutional credibility associated with a recognised European initiative.

In this sense, this initiative is not only a showcase for technology. It is also a practical instrument for helping innovative European SMEs gain market exposure, establish new business relationships and transform research and innovation results into commercial opportunities.

Voice from the participating organisations

The feedback received from the participating organizations was very positive, and the following three themes kept recurring: visibility, collaboration, and credibility.

Several organizations stressed the advantage of being part of a collective European presence rather than a separate one. INSIDE pavilion allowed not only interaction with external parties but also collaboration among the players our ecosystem itself.

As Diego Grimani from Innovation River noted:

“Following an internal debriefing, we concluded that for Innovation River,

Embedded World 2026 represented an unprecedented visibility opportunity. The event allowed us to meet many interesting people and organizations, and most importantly it gave us the chance to introduce and present our company to a wider community. We sincerely thank you for the opportunity offered and for creating such a collaborative and dynamic environment."

A similar perspective was shared by DAC.digital:
"Being part of the INSIDE Industry Association gave DAC.digital something that no trade show booth alone can provide: institutional context and the credibility that comes with it. Exhibiting at Embedded World 2026 under the INSIDE umbrella allowed us to show many years of engineering work across IoT, IIoT and multimodal agentic AI to the audience that understands it best."

This sentiment was echoed by VERUM, who highlighted both the professional and collaborative atmosphere of the pavilion. As they highlighted, "This year, the INSIDE booth was exceptionally versatile and full of familiar faces, which made it easy to spend the entire three days networking and sharing experiences within the booth itself. Embedded World resulted in several promising collaborations and interactions, and I can confidently say that these were three productive and well-spent days. The larger booth also attracted more attention to our stand, which was a great advantage. We are pleased to be part of INSIDE and grateful for the opportunity to showcase our company at Embedded World."

These testimonials demonstrate one of the key aspects of the initiative: apart from individual business opportunities, INSIDE pavilion provides an environment where companies can showcase their presence within a wider perspective of European excellence and collaboration.

Looking forward

The success of Embedded World 2026 has already laid the foundations for an even more ambitious participation in 2027.

The next edition, which will coincide with the 25th anniversary of the event, will take place from 16 to 18 March 2027 in Nuremberg. Preparations are already underway with plans to expand INSIDE pavilion from 125 m² to 180 m², increase the number of participating members to more than twenty-four, strengthen marketing activities and further enhance the visibility of European innovation on the global stage.

Particular attention will be devoted to broadening participation across the ecosystem, including stronger engagement with the European Design Platform, Design Enablement Teams (DETs), European projects and industrial stakeholders.

The objectives remain clear: strengthen European visibility, expand the ecosystem, create new business opportunities and accelerate the path from innovation to industrial impact.

Embedded World 2026 demonstrated that Europe possesses world-class technologies, world-class expertise and a vibrant innovation ecosystem. Embedded World 2027 will focus on making these strengths even more visible, more connected and closer to the market.

Because ultimately, technological excellence achieves its full value only when it leaves the laboratory and succeeds in the marketplace.

For more information flolow this link or scan the QR code

Participants

abinsula
DAC.digital INFOMOTION GROUP
SINETIQ
aitek
Innovation River
Cabildo de Gran Canaria
spegc
AITIA
ROPARDO
WHARD
Symbiotic EDA
cubit innovation labs
SIEMENS
verum.
D-CHIPLET
SAL SILICON AUSTRIA LABS
virtual vehicle
Chips JU

Engineering trustworthy industrial autonomy for Europe

AI, verification, and resilient logistics systems in the next generation of freight mobility

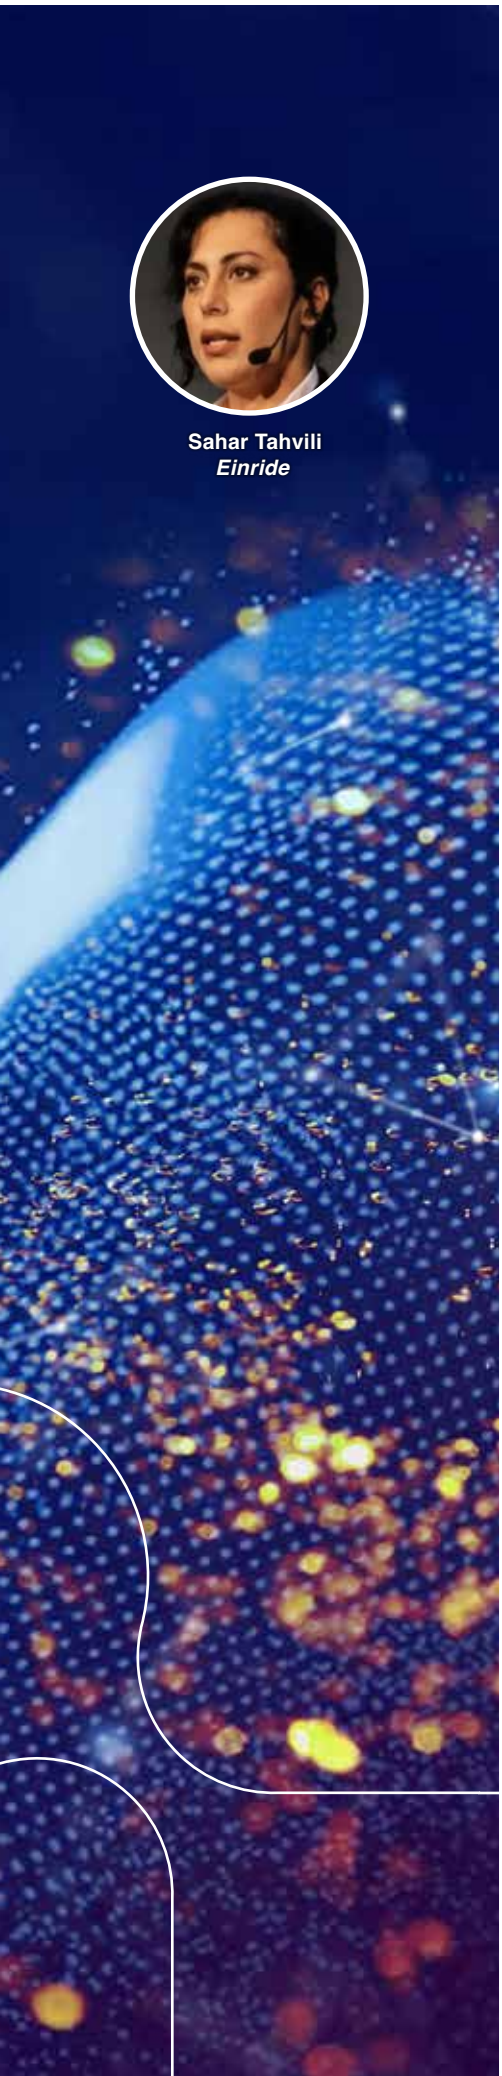

Sahar Tahvili
Einride

Europe is entering a decisive decade of industrial transformation driven by an unprecedented investment cycle in artificial intelligence, digital infrastructure, electrification, and autonomous systems. Across transportation, manufacturing, energy, and critical infrastructure, industries are under growing pressure to improve sustainability, resilience, safety, and competitiveness at the same time.

At the center of this transition is the convergence of artificial intelligence, automation, connectivity, and real-world industrial operations. These technologies are no longer evolving independently. They are becoming tightly integrated into complex cyber-physical systems that increasingly influence how industries operate, optimize resources, and make decisions. Freight transportation is one of the clearest examples of this shift.

Global supply chains are becoming more complex and less predictable. Energy transition targets, labor shortages, infrastructure limitations, and geopolitical uncertainty are placing new demands on logistics systems. At the same time, advances in AI, sensing technologies, cloud-native platforms, and edge computing are creating opportunities to rethink how freight networks are designed and operated. This transformation is often discussed in terms of electric vehicles or autonomous trucks. In reality, the change is much broader. What is emerging is a new generation of intelligent logistics ecosystems where vehicles, infrastructure, energy systems, operational platforms, and AI-driven decision-making work together continuously.

At Einride, autonomous and electric freight systems are being developed as part of an integrated digital ecosystem where transportation is orchestrated through software, operational intelligence, and real-time optimization. The lessons emerging from this transition extend beyond logistics itself. They provide insight into how Europe can engineer trustworthy industrial autonomy at scale.

From automation to intelligent logistics ecosystems

Historically, transportation innovation focused primarily on mechanical performance and infrastructure expansion. Today, logistics is increasingly becoming software-defined, connected, and data-driven. Modern freight operations generate large volumes of operational data related to traffic conditions, route planning, energy consumption, charging infrastructure, fleet utilization, warehouse coordination, weather conditions, and delivery timing constraints. Artificial intelligence enables these systems to move from reactive planning toward predictive and adaptive operations. In practice, this allows logistics networks to optimize routes, energy usage, fleet allocation, and operational efficiency continuously in real time. Electrification is a key part of this transformation. Electric freight systems support lower emissions and tighter integration with energy infrastructure. However, electrification alone cannot solve the growing operational complexity of modern logistics systems. AI-driven orchestration, automation, and autonomous capabilities are becoming equally important.

Autonomous freight systems are no longer limited to controlled research environments. Real-world deployments are increasingly demonstrating how autonomous operations can support industrial logistics use cases in structured and semi-structured environments. These deployments also reveal an important reality: scaling autonomy is not primarily a vehicle problem. It is a systems engineering problem. The challenge is not only building intelligent systems. It is building systems that can be trusted.

Trustworthy autonomy as a systems engineering challenge

In public discussions, artificial intelligence is often associated with computational power or model capability. In industrial environments, however, autonomy is defined by something much more demanding: the ability to operate

safely, predictably, and reliably under uncertainty. Industrial autonomous systems must function in environments shaped by changing weather conditions, infrastructure variability, sensor uncertainty, human interactions, and unpredictable operational events. They must also operate across regulatory boundaries, digital infrastructures, and increasingly interconnected industrial ecosystems.

This changes the engineering challenge fundamentally. Trustworthy industrial autonomy is not only about developing better AI models. It requires robust approaches to verification, validation, safety assurance, cybersecurity resilience, operational monitoring, and human oversight.

Autonomous systems must not only make decisions. They must provide confidence that those decisions remain safe and reliable under changing operational conditions. Traditional validation approaches alone are insufficient for AI-enabled systems because machine learning introduces probabilistic behavior and emergent operational characteristics. Exhaustively testing every possible real-world scenario is impossible. Instead, organizations increasingly rely on a combination of physical testing, large-scale simulation, digital twins, formal verification methods, and continuous operational monitoring to establish acceptable confidence levels for deployment. Engineering trust therefore becomes a continuous lifecycle activity rather than a final certification milestone.

At Einride, the development of autonomous freight technologies requires close integration between AI engineering, systems engineering, software development, operations, infrastructure integration, and safety validation. Trustworthy deployment depends not only on autonomous driving performance, but on the resilience and observability of the entire operational ecosystem surrounding it.

Verification, simulation, and operational AI assurance

One of the defining challenges of industrial autonomy is ensuring that intelligent systems remain trustworthy under highly dynamic and uncertain real-world conditions. As autonomous platforms become increasingly AI-driven, the engineering focus shifts from isolated model performance toward continuous operational assurance across the full system lifecycle. Real-world testing alone cannot practically cover the enormous diversity of operational scenarios, environmental conditions, infrastructure anomalies, and edge cases that autonomous systems may encounter. As a result, simulation, synthetic scenario generation, digital twins, and AI-assisted verification methodologies are becoming foundational components of modern autonomy engineering.

Advanced simulation environments now allow engineers to evaluate hazardous edge cases, degraded operational states, and rare-event behaviors at scales that would be impractical or unsafe to reproduce physically. AI is also increasingly used within the assurance process itself, supporting intelligent scenario generation, anomaly detection, adaptive testing, behavioral analysis, and operational risk assessment. This evolution is transforming verification and validation from a static certification activity into a continuous and data-driven engineering discipline. Trustworthy autonomy increasingly depends on the ability to combine:

- AI-driven scenario exploration
- Runtime observability and operational monitoring
- Data-centric validation pipelines
- Safety and resilience metrics
- Human-in-the-loop oversight
- Continuous operational feedback
- Traceability between system requirements and real-world behavior

Autonomous systems can no longer be treated as static deployments. They are adaptive cyber-physical ecosystems that continuously interact with infrastructure, operational environments, and human actors. Maintaining trust in such systems requires resilient engineering architectures capable of supporting transparency, explainability, safety, and continuous assurance throughout the operational lifecycle.

Similar challenges are already visible in other industrial domains. In telecommunications infrastructure, for example, AI-driven orchestration systems increasingly balance performance, resilience, and energy efficiency across highly distributed cloud-native environments operating under real-time constraints. These experiences demonstrate that trustworthy autonomy is not solely a machine learning challenge. It is fundamentally a systems engineering challenge requiring observability, adaptive assurance, continuous validation, and operational resilience at scale.

Europe's strategic opportunity

Europe already plays a central role in the global automotive and industrial landscape. The region has decades of experience developing safety-critical systems, advanced manufacturing capabilities, transportation infrastructure, and engineering standards that influence industries far beyond Europe itself. This industrial foundation creates a significant opportunity in the transition toward autonomous and AI-driven systems. At the same time, Europe is taking a leading role in shaping regulatory frameworks for trustworthy and human-centric AI. For safety-critical domains such as autonomous transportation, clear regulatory guidance and rigorous safety requirements are not obstacles to innovation, they are essential prerequisites for public trust and large-scale deployment.

However, Europe also faces an important balancing challenge. Industrial AI and autonomous systems are evolving at exceptional speed. Excessive regulatory fragmentation, long certification cycles, or unclear operational pathways can unintentionally slow experimentation, deployment, and industrial competitiveness. In fast-moving technology domains, innovation speed matters alongside safety and governance. The challenge for Europe is therefore not whether regulation is needed, but how to create frameworks that simultaneously support trust, interoperability, safety, and industrial agility.

This balance will likely become one of Europe's defining strategic questions over the coming decade. If Europe succeeds, it has the opportunity to lead globally not only in responsible AI governance, but also in the deployment of trustworthy industrial autonomy at scale. The combination of strong automotive heritage, industrial engineering expertise, digital infrastructure, and growing AI capabilities positions Europe uniquely to shape the future of autonomous industrial systems. In that context, trustworthy autonomy should not be viewed solely as a compliance challenge. It is also a competitiveness challenge, an infrastructure challenge, and ultimately a strategic industrial opportunity for Europe.

Industrial autonomy beyond transportation

The lessons emerging from autonomous freight systems are increasingly relevant across many industrial domains. Manufacturing, mining, energy, maritime systems, telecommunications, and industrial

robotics are all facing similar challenges involving AI-driven operational complexity, cyber-physical integration, safety assurance, scalable validation, and resilient human-machine collaboration. In this context, industrial autonomy should not be viewed simply as a mobility innovation. It represents the emergence of a new class of intelligent industrial infrastructure where AI-driven decision-making, connectivity, edge computing, distributed sensing, operational orchestration, and continuous assurance become deeply integrated into industrial operations.

The organizations and regions capable of engineering trustworthy versions of these systems will likely define the next generation of industrial competitiveness.

Engineering trust as Europe's competitive advantage

The global race in artificial intelligence is often framed around computational power, model sophistication, or the speed of innovation. In industrial environments, however, long-term competitiveness will depend less on who develops the most advanced AI models and more on who can deploy intelligent systems that can be trusted in real-world operations. This distinction is especially important in safety-critical sectors such as transportation, energy, manufacturing, and telecommunications. Autonomous systems in these environments must operate reliably under uncertainty while interacting continuously with physical infrastructure, human decision-makers, and changing operational conditions. Reliability, resilience, explainability, and operational accountability therefore become as important as algorithmic performance itself. Europe is uniquely positioned to lead in this transition.

The region combines strong industrial engineering traditions with a growing regulatory and societal focus on trustworthy AI. This creates an opportunity not only to develop advanced autonomous technologies, but also to define how these systems should be engineered, validated, governed, and integrated responsibly into society. The engineering focus is already shifting beyond isolated AI performance toward continuous assurance, operational observability, cybersecurity resilience, and transparent human oversight. Trustworthy industrial autonomy requires the integration of AI with systems engineering, safety engineering, resilient infrastructure, and operational governance.

The implications extend far beyond transportation alone. Industrial autonomy is becoming part of the foundational infrastructure of Europe's future industrial economy, influencing how goods are transported, factories are operated, energy systems are optimized, and critical services are coordinated. The transition toward intelligent industrial ecosystems is already underway. The challenge ahead is ensuring that autonomy is not only capable and efficient, but trustworthy enough to support long-term industrial and societal adoption.

In that context, trust is no longer simply a feature of autonomous systems. It becomes the enabling infrastructure for industrial autonomy itself.

Sahar Tahvili
Head of Verification and Validation at Einride and an adjunct associate professor of industrial AI systems at Mälardalen University.

With a background in applied mathematics and computer science, her work focuses on artificial intelligence, advanced software testing, autonomous systems, and decision support systems for safety-critical industrial environments. In 2022, she co-authored the book Artificial Intelligence Methods for Optimization of the Software Testing Process: With Practical Examples and Exercises, which was also recognized among the best new books in intelligent testing and AI-driven software quality engineering.

Technology Frontiers

The agent in the room

How autonomous AI agents are reshaping the threat landscape

Khalid Alnajjar
Senior AI & Data Scientist
F-Secure Oyj

Tuuli Lindroos
Cyber Policy Manager
F-Secure Oyj

Autonomous AI agents are moving from research prototype to consumer product at a pace that has outrun both the security tooling designed to protect users and the regulatory frameworks meant to govern them. Europe is regulating AI faster than almost anywhere in the world, the EU AI Act is in force, GDPR governs data, the Digital Markets Act is governing the structural conditions of platform competition and the Digital Services Act is reshaping platform accountability, and yet none of these frameworks were designed for software that browses, transacts, negotiates, and decides on your behalf without asking for permission at each step.

The security risks this creates are not simply extensions of existing threats. The most significant ones are structurally different, and they will not fully yield to the detection-and-blocking approach that has defined consumer cybersecurity for three decades. Some can be addressed. Some can only be managed. And some are properties of the technology itself that require fundamental research to address.

Agentic AI revolutionizing consumer cyber security

There is a scene already playing out in living rooms and on laptops across the world that would have looked like science fiction a decade ago. A person gives a brief instruction to a piece of software: “find me the best deal on a new car, my budget is €30,000.” Then they step away, perhaps spending their precious time with family and loved ones. The software browses listings, compares prices, fills in forms, initiates conversations with dealers’ automated systems, and eventually surfaces a recommendation. The person never sees most of what happened. They may not know which websites were visited, which services received their data, or which terms were implicitly accepted along the way. This is agentic AI. Software that doesn’t just respond to prompts but pursues goals, using tools and memory and judgment, across multiple steps, on your behalf. And it is arriving not as a gradual transition but as a sudden shift that the security industry, the regulatory community, and consumers themselves are still racing to keep up with.

For early adopters, the appeal is obvious. Install an agent, connect it to your accounts, and delegate the tedious work of modern digital life. What most users don’t consider, and what the broader industry has been slow to address, is that the same access that makes agents useful also makes them a target (a paradox unpacked by Paolo Azzoni in the previous issue of this magazine). The vulnerabilities are not the kind that a patch fixes.

From softbots to autonomous agents

Autonomous software agents are not new. The concept has been part of computer science since the early 1990s, when researchers began imagining “softbots,” software robots capable of carrying out

complex tasks on a user's behalf. A canonical early example was deceptively simple: tell the system you want to send a document to a specific person, describe roughly what it's about, and let the agent find the right recipient, locate the right file, and send it. Under the hood, this was a careful orchestration of logic trees and decision rules, pre-programmed paths along which the agent would navigate.

What made those early agents limited was

exactly what made them safe. They operated within narrow, well-defined domains. They couldn't reason about ambiguous situations, generalise to new contexts, or take initiative beyond what they had been explicitly programmed to do. Their decisions were traceable, their behaviour predictable.

The arrival of large language models (LLMs) changed this fundamentally. LLMs gave agents something they had never had before: the ability to reason across open-ended situations, communicate fluently in natural language, and pursue flexible goal-driven paths rather than rigid pre-programmed ones: the human just defines the goal and the agent reasons, performs actions and evaluates its way to completion.

In enterprise settings, agents have found their footing first. When a telecommunications operator deploys an agent to monitor network performance, detect faults, and trigger corrective actions, the goals are clear, the environment is controlled, and the outputs can be validated against measurable criteria. The same is true for agents monitoring transactions for fraud in banking, or for multi-agent fulfilment system, where specialized agents make continuous micro-decisions about routing, driver assignment, and order timing across a unified orchestration layer. These are bounded tasks, and bounded tasks are where agents perform well.

Personal assistants are an entirely different proposition. The use cases are open-ended, shaped by each individual's needs, habits, and tolerance for automation. One person might use an agent to manage their inbox and calendar. Another to track their children's school deadlines and announcements. Another to negotiate a car purchase. OpenClaw, a self-hosted personal assistant that connects to WhatsApp, Telegram, iMessage, and other messaging apps with full access to the user's filesystem and shell commands, browser, and external services and tools such as an email is a good illustration of where this appetite is heading: powerful, deeply integrated into daily life, and available to anyone willing to set it up. That pattern, an agent reached through a familiar channel, acting with broad permissions on the user's behalf, is now common across the category, whether open-source frameworks, domain-specific products like a shopping assistant, or major platform offerings like agentic browsers and general-purpose AI agents.

The architecture is also common across all of them, and it is worth understanding before we get to the threats. Every personal agent combines an LLM reasoning engine with access to the user's data and the ability to take actions in the world. The LLM at its core cannot reliably distinguish between instructions given by the user and instructions embedded in the data it accesses, a command hidden in an email looks the same as a command typed by the user. And because the agent operates on behalf of that user across tools, websites, and external services, it inherits this blind spot at scale. A tool, a website, another agent, or a malicious third party are all treated the same way: present, and assumed to be legitimate. There is no mechanism to question that assumption. This is not a bug. It is a structural feature of how these models process language, and it is the root of most of what follows.

The attack surface you can't see

When a consumer sets up a personal agent, they are not simply adopting new software. They are creating a new attack surface, one that extends from their own behaviour and trust, through the infrastructure hosting the agent, across every service the agent touches, and out into the open web. Understanding that surface requires looking at each component in turn.

The user

The first vulnerability is the human. Conversational interfaces trigger

interpersonal trust dynamics that evolved for human relationships, not software evaluation. Users disclose far more to agents than they would to equivalent non-conversational software: health concerns, finances, passwords, relationship problems. The same intimacy that makes agents feel useful also makes users less critical, more compliant, and more easily manipulated if the agent is ever compromised or spoofed. On top of this, users routinely grant agents far broader permissions than any single task requires. Full filesystem access, email access, payment credentials: handed over at setup, rarely reviewed again. This isn't carelessness. Broad access is what makes an agent useful in the first place. The value proposition and the attack surface are one and the same. And every permission granted amplifies the potential damage from every other vulnerability in the system.

The hosting environment

Open-source agents are typically self-hosted, running on personal laptops, home servers, or cloud instances configured by the user. Closed-source commercial agents shift that responsibility to the platform, but transparency varies widely: some openly share their security practices and agent behaviour in detail, while others offer only limited visibility into how user data is handled or what the agent does on the user's behalf. In both cases, the underlying risk is the same: sensitive credentials and personal data end up in infrastructure that was never hardened to protect them. For self-hosted agents, that means misconfigured ports, absent authentication, and API keys sitting unprotected in local files. For commercial platforms, it means trusting systems of varying opacity where a single breach can expose millions of users at once. These are not exotic attack scenarios. They are the predictable consequence of deploying powerful software without the guardrails that high-value data demands.

The tools and skills ecosystem

Modern agents extend their capabilities by connecting to plugins and skills, additional tools that allow them to perform specific tasks. This mirrors the extension and package ecosystems that have existed in software development for years, but with a critical difference: agent skills typically inherit the full permissions of the agent itself, meaning a single malicious skill can access everything the agent can. The marketplace where these skills are distributed is, at present, close to ungoverned. The ClawHavoc campaign, documented by security researchers in early 2026[1], found that around 12% of the skills on the ClawHub marketplace (for OpenClaw agents) were malicious: skills that appeared professional, carried fabricated positive reviews, and delivered infostealers and reverse shells running with the same OS-level permissions as the agent itself.

The underlying protocol for connecting agents to tools, the Model Context Protocol (MCP), has seen multiple high and critical CVEs since 2025, affecting both official and third-party components. One of these, CVE-2025-6514, affected mcp-remote, the OAuth proxy used by several major AI clients and code editors. It allowed full remote code execution when connecting to a malicious server, and affected nearly half a million downloads before a patch was issued. In another case, a widely used MCP email integration was found to have been trojanized, silently forwarding a copy of every outbound email to an attacker-

controlled domain. Every tool an agent connects to can receive sensitive data and act on the user's behalf, and a compromised tool can silently misuse both. A shopping tool connected to your personal agent could buy a birthday gift for your partner and one for the attacker too.

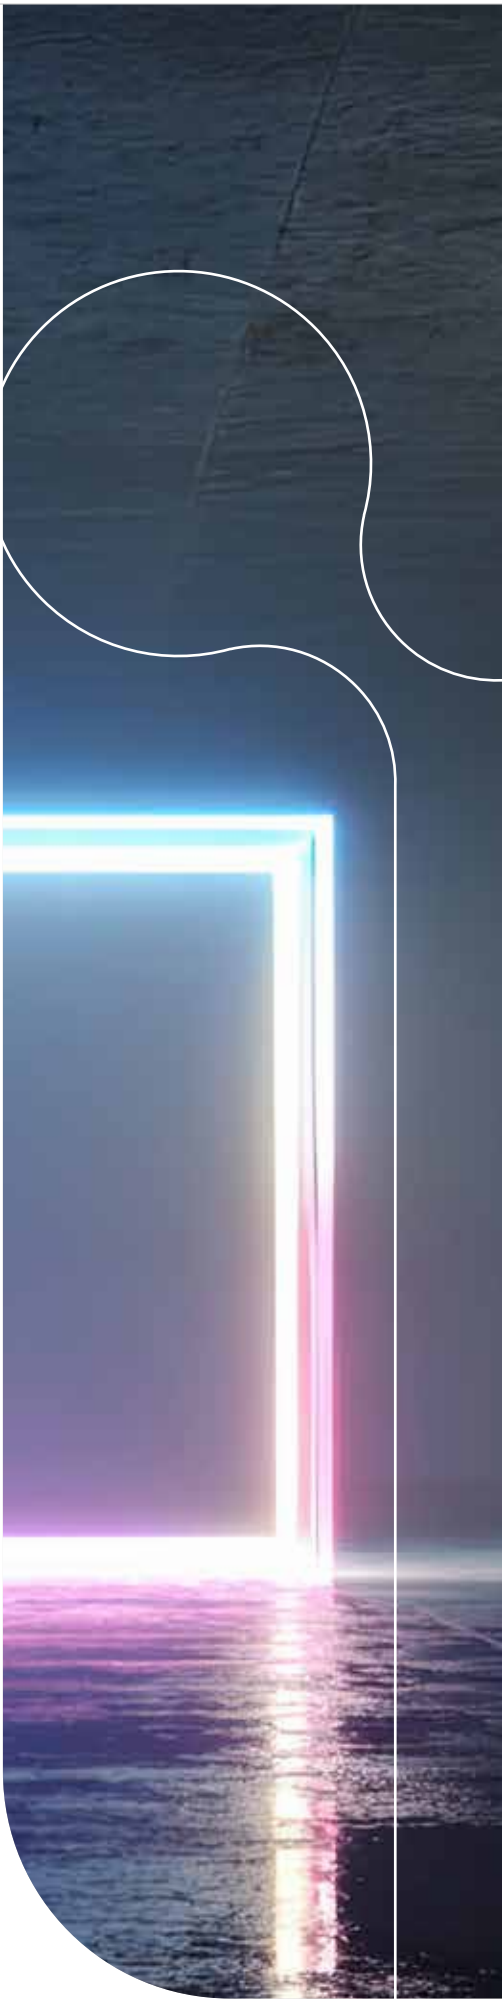

Key takeaways

- The security industry's traditional toolkit, detection rules, signatures, patches, is the wrong instinct for agentic AI. The threat landscape is not uniform, and frameworks that treat all threats the same will misallocate effort and fail consumers.
- A more useful framework divides threats into three tiers: those that can be addressed directly (malicious extensions, fraudulent apps, ungated agent actions); those that can be reduced but not eliminated (prompt injection, where action-level constraints matter more than language-level guardrails); and threats inherent to the current architecture, where no patch is the honest answer.
- Some risks are features, not bugs. LLMs are designed to be compelling and persuasive, measurably more so than humans. This creates consumer risk regardless of attacker intent, and calls for regulation and platform-level design rather than a security product.
- The next frontier is multi-agent ecosystems, where personal agents negotiate, book, and manage services autonomously on consumers' behalf. Each agent-to-agent interaction is a potential attack surface, and there is currently no equivalent of SSL certificates or domain reputation to establish trust in this environment.
- Building identity standards, reputation scoring, and anomaly detection for agent-to-agent interaction is one of the most consequential open problems in consumer security today.

At F-Secure, this is where our current research is focused, and where we are actively building the evidence base through international collaboration. Protection in a multi-agent world is not a single product problem, and it is not one any organisation can solve in isolation. It requires thinking seriously about how consumers can maintain meaningful control over agent ecosystems they cannot supervise directly. It requires risk-aware decision-making at the agent level, where the cost of over-blocking a legitimate interaction and the cost of missing a malicious one are both real and asymmetric. It requires trust verification infrastructure that does not yet exist. And it requires, perhaps most fundamentally, honesty about what security can and cannot deliver when some of the most significant risks emerge from the technology working exactly as designed.

F-Secure actively participates in co-innovation and research collaboration initiatives, to advance our work, maximise our impact within the tech community, and shape how protection is defined, designed and delivered. Through proactive engagement and outreach in initiatives like ELFMo, an ITEA4 Eureka Cluster project co-funded by Business Finland, and our Horizon Europe research partnerships spanning academia and industry across Europe, we are building a systematic understanding of the ecosystem layers and interaction dynamics that shape the multi-agent threat surface.

The LLM itself

A joint study across multiple AI safety institutions and labs published[2] in October 2025 demonstrated that as few as 250 carefully crafted malicious documents can successfully backdoor language models from 600 million to 13 billion parameters, and that the resulting compromise is undetectable by standard quality benchmarks. In practice, this means a small amount of poisoned training data could, for example, cause a model to consistently propagate medical misinformation while passing every standard accuracy test. This is a supply chain attack that reaches the consumer before they have installed anything. But the threat does not stop at training data. Agents that browse the open web face a growing volume of indirect prompt injections: malicious instructions hidden in web pages, designed to be invisible to humans but parsed and followed by AI systems. A recent scan of

the public web found a 32% increase in malicious prompt injection payloads between November 2025 and February 2026, while a separate investigation identified live payloads attempting financial fraud, data destruction, and credential theft on ordinary websites. The same malicious pages that exploit agents at runtime are also crawled into future training datasets, making the open web a vector for both immediate exploitation and long-term model corruption. The underlying problem is architectural: the model cannot distinguish between legitimate content and adversarial instructions embedded within it.

External interactions

When agents browse the web, they do so using the user's authenticated session, carrying their cookies, tokens, and credentials, and without the decades of anti-phishing defences built into conventional browsers or the protection layers offered by dedicated security products like F-Secure. A 2025 browser security study testing[3] over 100 real-world phishing attacks found that a leading browser blocked roughly half of them. An agentic AI browser blocked under 6%. Users of agentic browsers are approximately 90% more vulnerable to phishing attacks, not because the technology is careless, but because it was never built with this threat in mind.

And then there is what security researcher Simon Willison have come to call the Lethal Trifecta: the combination of access to private data, exposure to untrusted content, and the ability to take real-world actions. When these three conditions coexist, and in any capable personal agent they always do, a specific attack becomes possible. Malicious instructions embedded in content the agent encounters during a task, invisible to any human reader, can redirect the agent's behaviour entirely. This is indirect prompt injection, ranked by OWASP as the number one threat in large language model applications. OpenAI has publicly acknowledged that it is "unlikely to ever be fully solved."[4]

Conclusion: the agent is already acting

The personal agent represents something genuinely new in the history of consumer technology: software that acts as a delegate rather than a tool. It does not wait to be told what to do. It pursues goals, makes decisions, and takes actions in the world with your credentials, your data, and your authority. This is what makes agents valuable. It is also what makes them, in their current state, a significant and underappreciated risk.

The risk is not primarily one of malicious design. The major platforms building consumer agents are not trying to harm their users. The risk emerges from structural properties of the underlying technology, from deployment outrunning the security infrastructure needed to support it, and from a fundamental mismatch between the threat model agents create and the tools that currently exist to address it.

Europe has particular reason to pay attention. The EU AI Act's prohibition on subliminal manipulation and deceptive AI techniques became enforceable in February 2025. The Digital Services Act and its enforcement track record on dark patterns provides regulatory precedent. And the region's data protection framework, under GDPR, is increasingly strained by systems that aggregate, infer, and act on personal data in ways that existing consent and transparency mechanisms were never designed to address.

What comes next will be shaped less by any individual technology than by the choices made now: about what security obligations attach to agentic AI products, about what trust infrastructure gets built into the protocols connecting agents to services and to each other, and about how honestly the industry communicates to consumers what it can and cannot protect them from.

The agent in the room is already acting. The question is whether anyone is watching, and what watching even means when the actions happen at machine speed, across dozens of intermediaries, in conversations no human will ever read.

Khalid Alnajjar

Khalid Alnajjar holds a PhD in computer science, with an interdisciplinary research background focusing on multimodal NLP, machine learning, business and management, and cybersecurity. At F-Secure, he works at the intersection of AI and consumer protection, researching next-gen protections against scams and emerging digital threats, drawing on over a decade of experience bridging academic research and industry.

Tuuli Lindroos

Tuuli Lindroos is External Research Collaboration and Cyber Policy Manager at F-Secure's Futures Research Lab. Formerly at the Ministry for Foreign Affairs of Finland with experience in EU, UN, and OSCE technology negotiations, she monitors and analyses the cyber regulatory landscape and manages a multi-million euro portfolio of European research projects and partnerships.

1 https://www.koi.ai/blog/clawhavoc-341-malicious-clawedbot-skills-found-by-the-bot-they-were-targeting

2 https://www.anthropic.com/research/small-samples-poison

3 https://layerxsecurity.com/blog/layerx-identifies-vulnerability-in-new-chatgpt-atlas-browser/

4 https://openai.com/index/hardening-atlas-against-prompt-injection/

When you need to truly protect your data

An innovative off-grid approach to highly sensitive data with near-real-time access

Consider a personal bank account. On the surface, the data appears well protected; yet in practice it is accessible to current and former bank employees, domestic regulators, law enforcement agencies, and, in certain circumstances, authorities from foreign jurisdictions. For most personal banking, this level of exposure is acceptable. For high-value corporate secrets, critical infrastructure credentials, or state-sensitive information, it is not.

The same structural vulnerability applies to cloud storage. Under the US CLOUD Act (Clarifying Lawful Overseas Use of Data Act, 2018), US law enforcement may compel US-controlled cloud providers to hand over data regardless of where that data is physically stored, even if it resides on servers in the EU. In 2025, concerns intensified further when Microsoft publicly acknowledged that it "cannot guarantee data sovereignty" for EU customers should US government demands be issued[1]. The EU's own FISA Section 702, reauthorised in April 2024 with expanded scope, compounds the risk for non-US citizens whose data transits American infrastructure[2].

IoT devices that control sensitive infrastructure face an additional, distinct threat. Nation-state-backed attackers systematically target such devices: China's Volt Typhoon campaign maintained persistent access to US critical infrastructure by exploiting zero-day vulnerabilities in connected devices[3], and attacks on operational technology networks surged by more than 46% in Q1 2025 alone[4]. Compounding this, many device manufacturers retain the ability to push software updates remotely, meaning the supply chain itself (including the manufacturer and any cloud platform they rely upon) represents a potential access vector.

Regulatory note: The EU–US Data Privacy Framework does not resolve CLOUD Act exposure. Customer-controlled encryption with keys retained in-jurisdiction is the primary technical measure identified by the EDPB as capable of addressing this risk (EDPB Recommendations 01/2020).

Sensitive data deserves a vault

Physical analogy is instructive here: large sums of cash are kept in a vault, not on a desk. The same logic should apply to highly sensitive digital assets. Existing approaches each carry trade-offs:

- **Cloud vault services** (offered by major providers) provide convenience, but the provider and its governing jurisdiction retain theoretical access to your data, as the CLOUD Act discussion above illustrates.
- **Offline storage** (e.g. encrypted USB drives) eliminates network-based attack vectors but introduces inconvenience, requires physical security, and offers no near-real-time access for distributed teams.
- **Blockchain-based solutions** offer decentralised storage with strong integrity guarantees, but ownership of the cryptographic key remains the weakest link: the problem is displaced rather than solved.

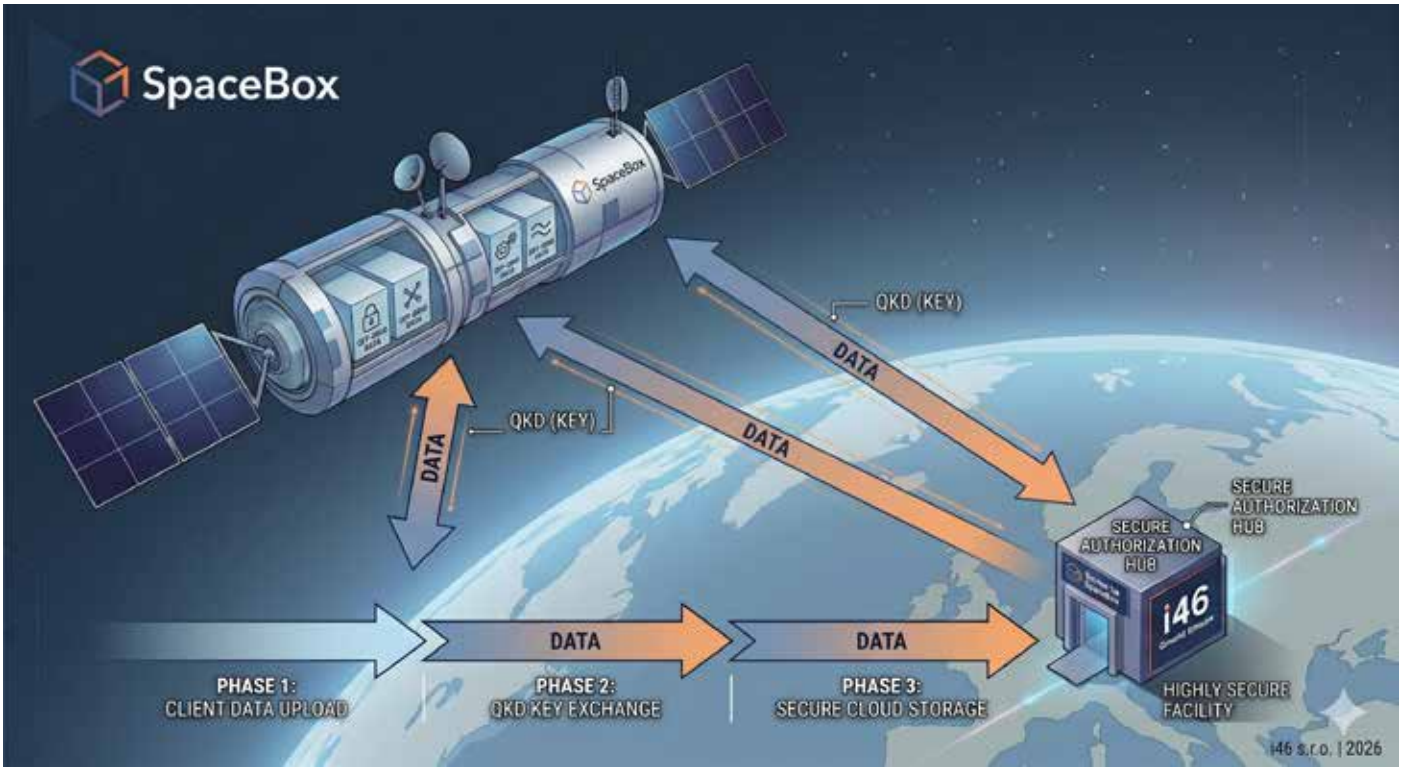

None of these approaches fully satisfies the requirement for a solution that is simultaneously air-gapped from terrestrial networks, physically inaccessible to adversaries, and still available on demand. This is the gap that SpaceBox addresses.

What is SpaceBox?

SpaceBox, developed by i46 s.r.o., provides organisations with their own off-grid digital storage in space. Two properties combine to deliver an unprecedented security profile: the off-grid architecture eliminates the conventional cyber-attack surface, and the physical location in orbit removes the risk of coercive or clandestine access through on-the-ground means.

Access to data stored in orbit is managed through a quantum-secured, rule-based authorisation system. Critically, the rules governing how data may be retrieved are defined by the customer before the data is loaded. This means that no party (including i46) can access customer data outside the parameters the customer has established. Examples of access policies include:

- Retrieval only upon approval by a quorum of designated custodians (e.g. three of five named executives).
- Cross-organisational authorisation, where a software update may only be deployed after sign-off by the device manufacturer, the technical support provider, and the end-user organisation simultaneously.
- Time-locked access, where data becomes available only after a predefined delay, giving all parties a window to identify and block unauthorised requests.

The quantum-secured channel ensures that the authorisation process itself cannot be intercepted or spoofed. Satellite-based quantum key distribution (QKD) is an active and rapidly growing field: the space-based QKD market is projected to grow from approximately $500 million in 2025 to over $1.1 billion by 2030[5], and Europe's EuroQCI programme is targeting a fully operational quantum satellite network within this decade[6].

Continuity without compromise: the dead man switch

One of the most underappreciated risks to organisational data security is not a cyberattack; it is the departure of a key executive. When a critical decision-maker leaves suddenly, sensitive credentials, authorisation keys, or classified project data can be lost or trapped in a personal access flow that no longer functions.

SpaceBox addresses this through an automated, time-based data transfer mechanism. If a designated custodian does not confirm continued custodianship within a predefined period (for example, 30 days), the system treats this as an inability to hold the data asset and automatically initiates a transfer to a pre-authorised successor. To prevent errors or manipulation, the transfer can be configured to require third-party countersignature before execution.

This mechanism provides a governance guarantee for business continuity: no single point of human failure can cause critical data to become inaccessible or uncontrolled.

Terrestrial SpaceBox: high assurance closer to home

Orbital storage offers unmatched security but comes at commensurate cost. Many organisations require strong data protection without the threat profile that justifies a space-based deployment. For these customers, i46 offers a terrestrial edition of SpaceBox, which the organisation hosts on its own premises. The data remains off-grid and inaccessible to external parties, while the access authorisation model (rule-based, multi-party, quantum-secured) remains identical to the orbital variant.

Current terrestrial deployments include:

- **Biometric data custody:** An organisation handling customer biometric records stores these off-grid in a terrestrial SpaceBox, satisfying data protection regulation while maintaining auditability and access control.
- **EV charger software update security:** A company responsible for updating the firmware of a network of electric vehicle chargers uses SpaceBox to ensure that

software updates are only deployed when the correct multi-party approval chain has been completed, preventing unauthorised or malicious firmware from being pushed to field devices.

- **Private 5G network key management:** A solution currently in preparation will store sensitive SIM card keys for a private 5G network off-grid, with keys made available to the network on authenticated request. This addresses a significant exposure point in private mobile network deployments, where SIM credentials are often stored in conventional, internet-connected systems.

The urgency of acting now

The threat environment is not static. IoT malware attacks surged by approximately 107% in 202[7], and nation-state actors are intensifying their focus on critical infrastructure worldwide. Regulatory pressure is also increasing: NIS2 (applicable from October 2024) requires essential and important entities to assess supply chain ICT risk, including provider exposure to non-EU government access demands (a direct reference to CLOUD Act risk)8. DORA (applicable from January 2025) requires financial entities to evaluate concentration risk in their ICT relationships[2]

.

Organisations that delay addressing data sovereignty will find their options progressively narrower as both threat actors and regulators raise the bar.

i46 as your data sovereignty partner

i46 s.r.o. is not a one-size-fits-all provider. Every deployment of SpaceBox, whether orbital or terrestrial, is configured to the specific data classification, access policy, and governance requirements of the customer organisation. The design principle is that the customer retains full sovereignty over their data: i46 provides the platform and the assurance, not the keys.

We welcome the opportunity to understand your organisation's data security challenges and to design a solution that is proportionate to the risks you face.

About i46 s.r.o

i46 s.r.o. provides cybersecurity solutions with a specialised focus on data sovereignty and regulatory compliance. The company's solutions include network analysis tooling, encryption services for devices without secure elements, SBOM analysis, and the SpaceBox platform for off-grid secure storage. i46 is a partner in Moore4Power, a CHIP-JU project. Spacebox by i46 s.r.o., has been supported by the ESA Technology Broker Czech Republic through the ESA Spark Funding initiative.

1 Claromentis. "Understanding the implications and risks of the US CLOUD Act" (2025). Discusses Microsoft's 2025 admission regarding EU data sovereignty. claromentis.com

2 CMS Law. "Demystifying the debate on the US CLOUD Act vs European/UK Data Sovereignty" (February 2026). Covers CLOUD Act mechanics, FISA 702 reauthorisation and NIS2/DORA obligations. cms-lawnow.com

3 Waterfall Security Solutions. "Learning From 2024's Top OT Attacks and Planning for 2025's Security." Details Volt Typhoon and nation-state escalation against critical infrastructure. waterfall-security.com

4 DeepStrike. "IoT Hacking Statistics 2025: Threats, Risks & Regulations." Reports 46% increase in OT ransomware Q1 2025 and 820K daily IoT attacks. deepstrike.io

5 The Quantum Insider. "Space-Based Quantum Key Distribution: Market Map and Competitive Landscape 2025" (March 2025). Market sizing and adoption timeline. thequantuminsider.com

6 TS2 Space. "Quantum Leap: Satellite QKD's Race to Secure the Global Data Economy (2024–2031)." Covers EuroQCI / EAGLE-1 mission and global QKD programmes. ts2.tech

7 CM Alliance. "Top 10 Biggest Cyber Attacks of 2024." Cites SonicWall Cyber Threat Report: 107% surge in IoT malware attacks in 2024. cm-alliance.com

8 ISACA. "Cloud Data Sovereignty: Governance and Risk Implications of Cross-Border Cloud Storage" (2024). Covers NIS2 supply chain obligations and GDPR jurisdictional risk. isaca.org

Innovation River at Embedded World 2026

Opening new pathways through collaboration

Diego Grimani
Innovation River

Participation in Embedded World 2026 as a co-exhibitor alongside the INSIDE Industry Association marked a significant milestone in Innovation River's growth trajectory.

Held in Nuremberg, Embedded World stands as a global reference point for the embedded community, bringing together leading experts, industry stakeholders, and research organizations from across the world. The event provides a unique platform to explore the full spectrum of embedded technologies, from hardware and software to AI, communication systems, and complex system design, while fostering dialogue on future trends and emerging challenges.

For Innovation River, this represented the first opportunity to take part in such a high-level international exhibition. Thanks to the valuable support of INSIDE, we were able to immerse ourselves in a dynamic and highly innovative environment, where knowledge exchange, networking, and collaboration are central pillars. Embedded World is not merely an exhibition, but a vibrant knowledge hub where top experts share insights and forward-looking perspectives through conferences, panels, and in-depth technical discussions.

Embedded World 2026 was both highly valuable and profoundly transformative. We significantly expanded our professional network, engaged with key stakeholders, and showcased our activities to a distinguished international audience. Despite being a growing company, Innovation River successfully positioned itself within a competitive and forward-looking ecosystem, gaining increased visibility and recognition.

Most importantly, several concrete opportunities have already emerged from the connections established during the event, further confirming the strategic importance of participating in such a prominent global platform. The quality of interactions and the openness of the community clearly demonstrate why Embedded World is widely regarded as a must-attend event for the industry.

We would like to express our sincere appreciation to the INSIDE Industry Association for enabling our participation as co-exhibitors and for actively fostering collaboration across the European innovation landscape.

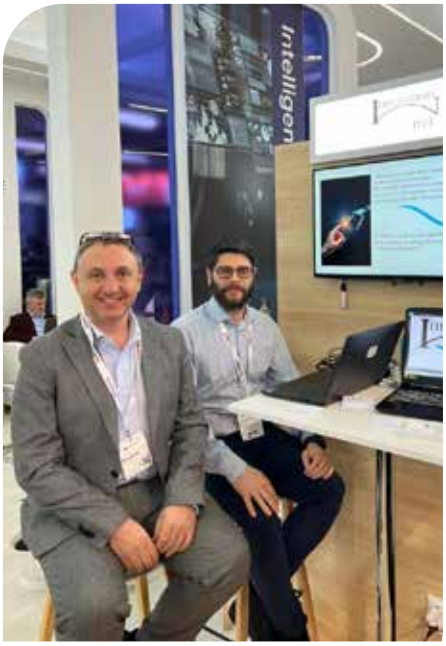

At Innovation River, we move forward with renewed enthusiasm, strengthened partnerships, and a clear vision for the future, building on the momentum generated at Embedded World 2026 and proudly participate in the 2027 edition, reaffirming our commitment to advancing the embedded ecosystem.

CHIPS ACT 2.0

A major step forward for Europe's technological sovereignty

The European Commission has unveiled the proposed **Chips Act 2.0**, as part of the broader Technology Sovereignty Package.

At a time when semiconductors power everything from AI and mobility to healthcare, defence and energy systems, Europe is reinforcing its ambition to strengthen its semiconductor ecosystem and reduce strategic dependencies.

The proposal focuses on:

- Strengthening research, innovation and skills
- Accelerating strategic investments
- Enhancing Europe's resilience in critical technologies
- Supporting stronger links between chip manufacturers and end-user industries

As the voice of Europe's Electronic Components and Systems community, **INSIDE** welcomes this important step towards a stronger, more competitive and more resilient European ecosystem of electronic components and systems.

€1.37 TRILLION

Global semiconductor market by 2030

AI

AI-related apphications expected to drive around **70%** of future growth

Powering the technologies that will shape our future:
AI, mobility, healthcare, energy, industry automation, defence, etc.

A tremendous opportunity and a strategic challenge for Europe.

The Chips Act 2.0:
A framework to support innovation, investment and industrial growth across Europe.

At **INSIDE**, we are committed to supporting the Electronic Components and Systems community in shaping a competitive and resilient Europe.

NO CHIPS WITHOUT SKILLS

Europe's electronic components and systems ambitions depend on people.

The Chips Act 2.0 highlights the importance of developing the skills, talent and expertise needed across the entire semiconductor value chain.

ATTRACT TALENT
Inspiring the next generation to build the technologies of the future.

DEVELOP SKILLS
Equipping people with the right skills for today and tomorrow.

BUILD EXPERTISE
Strengthening technical and scientific expertise across the entire value chain.

STRENGTHEN COMPETITIVENESS
A skilled Europe is a competitive and resilient Europe.

At INSIDE, we are committed to supporting the Electronic Components and Systems community and to invest in Europe's greatest asset: its people.

A STRONG ECOSYSTEM A COMPETITIVE EUROPE

The Chips Act 2.The Chips Act 2.0 highlights the importance of collaboration across the entire electronic component and systems value chain.

Together, we innovate.
Together, we build Europe's future.

RESEARCH
Pushing the boundaries of innovation

DESIGN
Building advanced, intelligent and energy-efficient solutions

MANUFACTURING
Scaling up excellence in Europe

PACKAGING & TESTING
Enabling performance and reliability

APPLICATIONS
Powering Europe's industries and society

INSIDE connects and represents Europe's Electronic Components and Systems community.

Together, we build innovation, resilience and strategic autonomy.

FROM DEMAND TO SUPPLY

Chips Act 2.0

The next challenge is not just producing chips.

EU needs ECS that satisfy the application requirements:
EU's strength lies in application

STRATEGIC APPLICATION AREAS

- Automotive & Mobility
- Healthcare & Medical Technologies
- Industrial Automation
- Energy Systems
- AI Infrastructure
- Aerospace & Defense

VALUE-CHAIN COVERAGE AS A STRATEGIC PRINCIPLE

- Move beyond isolated technology investments
- Evolve together across the full value chain
- Create European ecosystems that compete globally from application to chip

- Intelligent systems powered by software and data
- Software creates value and competitive advantage
- ECS are driven by the demand generated by applications

TECHNOLOGY FOLLOWS DEMAND.

APPLICATIONS → SYSTEMS → SOFTWARE → ECS

APPLICATION-DRIVEN INNOVATION FOR EUROPE'S TECHNOLOGICAL SOVEREIGNTY

From integration complexity to integration agility

Karl-Johan Gramner
Sinetiq

Why API governance and interface validation are becoming strategic capabilities for resilient digital systems

As digital systems become increasingly interconnected, integration is no longer a secondary technical concern. It has become a strategic capability. Across industry, public services, and research-driven environments, organisations depend on complex landscapes of applications, services, platforms, and data flows. Yet in many cases, the ability to integrate these systems remains fragile, slow, and difficult to govern.

This challenge is particularly visible in organisations operating across multiple teams, suppliers, or legacy environments. APIs are often developed in parallel, without cross boundary synchronization, resulting in limited alignment and reuse. Interface issues are discovered late, sometimes only when systems are already being connected. Visibility into deployed services, ownership, versions, and runtime behaviour is frequently incomplete. The result is familiar: duplicated work, delayed delivery, inconsistent quality, and hidden operational risk.

For Europe's digital and industrial ecosystem, this is not a marginal issue. Interoperability, resilience, and controlled yet flexible futureproof system evolution are increasingly central to competitiveness. As digital infrastructures expand across sectors, the quality of integration becomes a determining factor in whether innovation can scale reliably.

Integration problems are often organisational problems in technical form

When integration fails, the symptoms usually appear technical. Interfaces do not match. Documentation is incomplete. Dependencies are unclear. Deployment reveals incompatibilities that were not visible in development. Yet these visible issues often reflect deeper structural challenges. A common problem is lack of visibility. Teams may not know which APIs already exist, who owns them, which version is active, or how services are being used in practice. In such environments, reuse becomes difficult and duplication becomes almost inevitable.

Timing is another critical factor. When interface validation happens late in the development process, small inconsistencies can trigger costly cycles of rework and coordination. By the time systems reach integration testing, the cost of correction is already high. A third issue is the growing difficulty of change. Systems are expected to evolve continuously, but their integration models are often too opaque or too rigid to support controlled adaptation. This becomes especially problematic in environments where security, traceability, and operational continuity must coexist.

Seen from this perspective, integration is not simply about connecting systems successfully. It is about creating the conditions under which systems can evolve with confidence.

Why governance matters more as complexity grows

API governance is sometimes perceived as a constraint, an additional layer of control imposed on development teams. In practice, effective governance does the opposite. It reduces uncertainty, improves visibility, and allows teams to move faster with fewer surprises. In complex environments, governance becomes valuable when it connects design intent, ownership, lifecycle, and implementation. API specifications, deployed services, and runtime behaviour are too often treated as separate concerns. When they remain disconnected, organisations lose the ability to understand what was designed, what was implemented, and what is actually running.

A governance-oriented approach helps restore that continuity. It creates a reliable view of active interfaces and their evolution over time, while supporting stronger traceability. This makes it easier to manage versions, reduce duplication, and align delivery with architectural intent. In large organisations and distributed ecosystems, this becomes essential. Interoperability depends not only on technical quality, but also on shared understanding. Governance, in this sense, is not a bureaucratic layer, it is a practical enabler of coherence.

Interface validation shifts quality upstream

If governance provides visibility, interface validation provides confidence. The earlier organisations can detect mismatches, the lower the cost of correction and the lower the risk of cascading delays. Interface-first approaches play a key role here. By validating APIs independently of full system integration, teams can detect and resolve issues within their own environments before they become cross-organisational problems. Provider and consumer behaviour can be simulated, dependencies can be mocked, and assumptions can be tested earlier in the lifecycle.

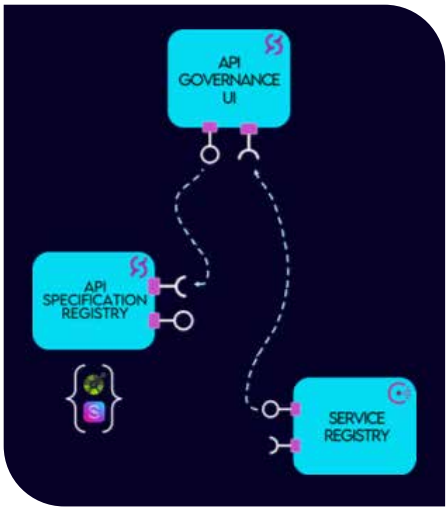

This fundamentally changes development dynamics. It reduces the number of integration cycles, shortens the path to production, and improves predictability across teams. Instead of discovering issues when multiple systems are already coupled, organisations can surface them earlier, when they are still manageable. In environments with many interfaces and evolving service landscapes, early validation also reduces friction between teams. It supports a more modular delivery model, where progress does not depend entirely on the readiness of every connected system. The value is not only speed. It is the ability to improve delivery quality without relying on late-stage troubleshooting to reveal structural weaknesses.

From isolated fixes to integration agility

A useful way to frame this challenge is through the concept of integration agility. The key question is not only whether systems can be connected, but whether they can adapt to change in a controlled and sustainable way. This requires a broader perspective than API design alone. It includes the ability to inspect systems, observe behaviour, secure interactions, manage dependencies, and respond to failures with minimal friction. Integration quality depends on a wider set of system characteristics, including composability, elasticity, inspectability, observability, resilience, and security.

These dimensions matter because integration is never static. Systems evolve, services are replaced, interfaces change, and operational contexts shift. Organisations that manage this effectively are not necessarily those with the most advanced technologies, but those with the clearest structures for handling change. This perspective is particularly relevant in Europe's current digital landscape, where interoperability, resilience, and technological sovereignty are increasingly interconnected. Systems that can reliably share services and data are better positioned to support both innovation and long-term stability.

Making system landscapes visible again

One of the recurring challenges in modern IT environments is the gradual loss of visibility. Documentation drifts, runtime reality diverges from architectural models, and ownership becomes fragmented. Over time, the system landscape becomes harder to understand and manage. Restoring visibility is therefore not simply an operational improvement, it is a prerequisite for better decision-making. Organisations need to know what exists, what is active, what is reused, and where risks are concentrated.

Approaches that combine specification registries, service registries, and governance interfaces are particularly valuable in this context. They help create continuity across the lifecycle from contract to implementation to runtime enabling more informed decisions and reducing hidden complexity. The same applies to testing and validation environments. When interfaces can be simulated and verified early, teams gain a clearer understanding of how change will propagate. This strengthens both delivery confidence and operational readiness.

A strategic capability, not just a tooling issue

Ultimately, integration quality is not just a tooling concern. It is an organisational capability with strategic implications. When governance is weak and validation is delayed, complexity accumulates silently. Delivery slows, costs increase, and risks become harder to identify. By contrast, organisations that treat API governance and interface validation as integral parts of system architecture create better conditions for speed, reuse, and resilience. This does not eliminate complexity but it makes it more visible, more structured, and more manageable.

For Europe's innovation ecosystem, this distinction is becoming increasingly important. As digital systems grow more interconnected across sectors and borders, the ability to govern interfaces and validate integration early will play a key role in transforming technological potential into reliable deployment.

Conclusion

Digital transformation is often framed in terms of platforms, data, artificial intelligence, or cloud infrastructure. Yet many of the practical barriers to progress emerge at the level of integration where systems meet, where interfaces fail, and where organisational complexity becomes technical friction. API governance and interface validation therefore deserve greater strategic attention. They are not peripheral technical functions, but essential components of the infrastructure that allows digital systems to evolve safely, predictably, and at scale.

For organisations navigating increasingly complex system landscapes, the challenge is no longer simply to connect technologies. It is to do so in a way that preserves visibility, supports change, and reduces uncertainty over time. In this context, integration agility is not just a technical objective, it is a condition for resilient innovation.

This perspective reflects the growing importance of structured integration practices, an area in which Sinetiq are actively contributing through governance and validation-focused approaches, including tools, methods and supporting consultancy services. Specifically, the abilities described above are provided by the Sinetiq products API Governance Tool (AGT), Extensible Test Tool (xTT) and Integration Agility Assessment (IAA).

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

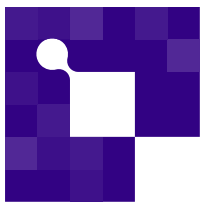

INSIDE
Industry Association

INSIDE-association.eu