

EURAIKA

Trustworthy AI. European engineering.

THE MANIFESTO

Why we are building a sovereign AI stack for Europe — and how

Bert Colemont

Founder & Visioneer, Euraika

A first-person statement of what Euraika is for. It sits alongside the Euraika Brand Book and the 2026 product portfolio. Where claims are quantified or personal, confirm with the team before anything goes public.

I built this on purpose

I started Euraika because I could see where this was going. Artificial intelligence is about to change how every organisation in Europe works — how a hospital handles a record, how a bank clears a payment, how a public administration serves a citizen. And I could see, just as clearly, that Europe would be right to demand this be done responsibly. We have written that demand into law: GDPR, NIS2, DORA, the EU AI Act. Two waves were going to meet — AI's leap forward and a hard new European compliance reality — and I wanted to be standing exactly where they collided.

Most of the market made a different choice. It chose convenience. It said: adopt AI, but send your data abroad, accept a model you cannot inspect, and surrender control of the system you now depend on. I refused that trade. We chose to do it properly. Everything that follows is a consequence of that one decision.

We never trade your sovereignty for convenience — and we refuse to make sovereignty feel like a burden.

The name is an argument

Euraika is not a coincidence of letters. Eur + AI + eureka — Europe, artificial intelligence, and the cry attributed to Archimedes of Syracuse at the moment a hard problem suddenly resolves. The whole promise is folded into the word: European AI, and the clarity of a problem finally solved the right way.

I wear that lineage openly because it tells you how we work. Archimedes did not chase magic. He reasoned from first principles, he built the machine, and he showed his working. That is the spirit Euraika is built in: engineering over shortcuts, capability over claims, proof over promises. It is also why we refuse the easy trade the rest of the market accepts. I would rather solve sovereignty and usability properly than bolt either one on at the end.

The parts of our stack that deliver, protect and guard carry names from Greek mythology — a small, coherent pantheon. That is deliberate too. A name should carry meaning, not noise.

How we began

Euraika started as a conversation. Alicia Fawcett and I got talking about AI and where it was heading, and decided to build a company. We brought in Alicia's husband, Jorge, and a Euroclear colleague, Frederic Collard. Our first ambition was enormous and a little naive: to become the European alternative to ChatGPT.

The team grew. I invited Irene Personne, and — because I never do business without him — my accountant Kris Dewalheyns came on board too. Then a concrete need changed our course: the chance to build a real application for the Belgian federal pension department. That need pointed us somewhere sharper and more European than a general chatbot. It pointed us at trustworthy, compliant, sovereign software.

Around that pivot the founding team settled. Alicia and Jorge chose to step away, and we continued as four founders — Frederic, Irene, Kris and myself — together with the engineers

who built our first product, Aegis: Andrei, Laurence, Amin, Khaula, Paul and Yvan. A fifth founder, Thomas Baeck, then joined. Thomas understands the market and how these tools should really work, and with him we set our course toward a full sovereign security stack, with Aegis as the starting point.

The vision is mine; the company is deliberately shared. I am Euraika's visioneer — the source of the ideas and the product thinking — but I never wanted to build alone. I know my own strengths and my own limits, so I chose co-founders whose talents complete the picture. Every founder holds an equal share and an equal vote. No single person owns the truth here. A partnership of equals, by design.

What we are for

Let me say it plainly, in my own words.

I started Euraika because I could see where this was going. AI would change how we work, and Europe would rightly demand that it be done responsibly. Most of the market chose convenience. We chose to do it properly.

I am the visioneer, but I never wanted to build alone. I know my strengths and my limits, so I chose co-founders whose talents complete the picture — and we hold equal shares and an equal vote. No single person owns the truth here.

We build software that respects the people who use it and the continent we build it in: sovereign, accountable, and a pleasure to use. Proof, not promises.

My vision is bigger than any single product. I want Euraika to be the best place in Europe to work — a flat company where everyone is equal and psychologically safe — and the top European AI software company: sovereign software at a democratic price, within reach of every organisation, not just the largest. That is how we win Europe — by being trusted, and by being for everyone.

The vision has two halves, and they hold each other up. Inward: the best place in Europe to work — flat, equal, psychologically safe, where people do their best work without fear. Outward: the top European AI software company, winning the market with sovereign software priced so organisations of every size can afford it, not just the biggest. A healthy, equal team is precisely what makes software that people love. You cannot fake one without the other.

The thesis: European digital sovereignty

This is our centre of gravity. Europe runs its most sensitive operations — hospitals, banks, public administrations, critical infrastructure — increasingly on AI. Yet almost all of the AI on offer asks the same price: export your data, trust a black box, and give up control of the system you depend on. For a continent with GDPR, NIS2, DORA and the EU AI Act written into law, that price is no longer acceptable — not to regulators, not to boards, not to citizens.

My position is simple. Sovereignty is not a feature you bolt on. It is an architectural decision you make at the start. We build AI software the European way — deployed on infrastructure the customer controls, governed transparently, accountable by design — so European organisations can adopt AI without giving up control of their data, their models, or their compliance posture.

Why now

Three forces are converging, and Euraiika lives at their intersection.

- **Regulation has teeth.** NIS2 and DORA are in force; the EU AI Act phases in through 2026–2027. Compliance is no longer slideware — it demands continuous, auditable evidence.
- **Sovereignty is a board-level word.** Cloud concentration risk, geopolitical exposure and data-transfer uncertainty have moved “where does our data live?” out of the IT department and into the boardroom.
- **AI trust is the bottleneck.** Organisations want AI’s productivity but cannot deploy black boxes into regulated workflows. The blocker is governance, not capability.

The antagonist I name is not a competitor. It is the status-quo trade-off itself: “adopt AI, but only by exporting your data and trusting a black box.” We remove that trade-off. That is the story — and I tell it constructively, never with alarm.

The second pillar: effortless by design

Sovereignty is what we protect. Simplicity is how we win the person actually using the software. These two pillars carry equal weight in everything we do.

Most enterprise compliance, security and AI-infrastructure software is powerful and punishing — built for specialists, gated behind training, manuals and consultants. I reject that completely. We build to the standard I hold in my head as the Apple standard: a product so intuitive that any competent professional can sit down, without a manual, and accomplish something real on the first try. Simple on the surface, powerful underneath.

Here is what that means concretely, and what we can honestly claim:

- **No manual required.** If a user needs documentation to complete a core task, the design has failed. The interface teaches itself.
- **Simple, not simplistic.** We hide complexity, we don’t remove capability — progressive disclosure, not dumbing down.
- **It just works.** Sensible defaults, opinionated paths, one obvious next step. The software makes the right thing the easy thing.
- **Calm, uncluttered.** Visual restraint matches our tone of voice: composed, precise, unembellished. The interface never shouts.
- **Fast to first value.** A new user reaches a meaningful outcome in minutes — a compliance status, a deployed control — not after a six-week implementation.

Our Brand Book says “Sovereignty over convenience,” and people sometimes misread it. It is about trust trade-offs: we will never weaken sovereignty, governance or accountability to make something more convenient. It is not a licence for clumsy software. Inside the

boundaries of sovereignty and trust, the experience must be effortless. Sovereign, compliant software has a deserved reputation for being clunky. We make the responsible choice also the easy one. “Powerful, compliant, sovereign — and a pleasure to use” is a position almost no one in this category can credibly claim. We can.

Who we are

These are our canonical, company-level statements, and I stand behind every one.

- **Category:** European AI software development.
- **Mission:** Euraika develops trustworthy AI software that European organisations can deploy, govern and rely on — with full control over their models, their data and their infrastructure.
- **Vision:** A Europe where organisations lead in sovereign AI — deploying, governing and trusting their own intelligent systems by design.
- **Brand promise:** Sovereignty over convenience. Capability over claims. Engineering over shortcuts.
- **Anchor:** Euraika builds AI software that European organisations can deploy on their own infrastructure, govern transparently and trust with evidence — not assumptions.

And what we are not: not a consultancy, not a black-box vendor, not a US platform with an EU data-residency checkbox. We are engineers and builders.

Five values survive into everything we make: Accountability by Design, Transparency through Evidence, Technological Sovereignty, Engineering Excellence, and Ethical Innovation. AI assists; humans decide.

The heart of the company

Our ideas are born in R&D, and I lead it myself. It is the creative engine where products take shape long before they reach a roadmap. It is also, deliberately, a neurodivergent team — minds that are ADHD and/or autistic — and I treat that as a core creative strength, not a footnote. Pattern-finding, deep focus, lateral thinking, and an honest, low tolerance for anything needlessly complicated are exactly the instincts that produce software which is powerful underneath and effortless on the surface. The way we work is the reason the products feel the way they do.

Every application we build has to answer “yes” to four questions:

- **Is it easy to use?** Usable without a manual, intuitive from the first try.
- **Does it improve the user’s quality of work and life?** It removes friction and gives time back — it does not add busywork.
- **Is it built from the daily user’s perspective?** Designed around the real person doing the real job, not an abstract buyer or a feature list.
- **Does it create value for Europe?** Sovereign, accountable, and worth building on this continent — not just technically possible.

The portfolio is one story

Euraika is not a collection of unrelated tools. It is a layered sovereign AI stack. Each part owns one layer, and they share one set of principles: EU-hosted, self-hostable, governed, evidence-producing, human-in-the-loop.

Read it from the bottom up. Hermes is the engine underneath — not something we sell, but the EU-hosted inference that carries every app above it. On that foundation, a European organisation proves it is compliant (Aegis), watches the threat landscape (Cerberos), and runs its business (RunBiz) — all without a single byte leaving European, customer-controlled infrastructure.

There is a thread I lean on hard, because it is true: Aegis is not only a product, it is also the governance discipline inside Hermes. The same evidence-and-verification rigour that powers our compliance command center is what makes our AI inference trustworthy. We eat our own cooking.

Two product lines run on the same foundation. The security stack, led by Thomas Baeck, starts with Aegis and extends through Cerberos (threat intelligence), Athena (phishing and security awareness), and an AI-driven automated-pentest product still to be named. The arc: know your compliance posture, watch the threats, train your people, continuously test your defences — all sovereign. Separately, RunBiz is our EU-first CRM-to-ERP line for everyday business work. Both run on Hermes. Project Camiel sits outside both, as an open-source community project.

Hermes — our sovereign inference engine (not for sale)

Hermes is a Kubernetes-native inference backend that serves open-source and EU-hosted models to our own applications, behind OpenAI- and Anthropic-compatible APIs, with the Aegis governance discipline built in — retrieval, claim verification, citations. It is the reason I can promise that no data leaves European, customer-controlled infrastructure. It is the literal proof that the governance we sell in Aegis runs on every model call our apps make. I treat Hermes as a credibility asset, not a pitch.

In one line: Hermes is the EU-hosted engine inside every Euraika app — the reason our sovereignty is real, not a label.

Aegis — AI-powered compliance command center (our flagship, ready today)

Aegis turns governance, risk and compliance from periodic panic into a daily operation. It produces auditable evidence for European frameworks — GDPR, NIS2, ISO 27001, DORA, CyberFundamentals, the EU AI Act — and keeps a human accountable for every decision. It is built for Data Protection Officers, CISOs, compliance leads and IT directors in regulated industries. Single-tenant deployment, a full audit trail on every action, a real evidence architecture, AI features clearly gated and governed. We compete with OneTrust, Vanta, Drata and Secureframe — but we are built from European regulation, not retrofitted from US frameworks.

In one line: Aegis turns compliance from periodic panic into continuous, evidence-backed operation.

Cerberos — OSINT threat intelligence (alpha)

Cerberos is a threat-intelligence platform focused on the EU and Belgium that collects, normalises, enriches, scores and serves intelligence from public and open sources — designed to work without paid API keys. Threat intelligence today is expensive, US-centric and locked behind proprietary feeds, and European context is thin. Cerberos gives analysts and SOC teams a European view of the threat landscape, on their terms.

In one line: Cerberos watches the European threat landscape — from open sources, without the proprietary-feed tax.

Athena — phishing & security awareness (roadmap)

Athena will run realistic, safe phishing exercises, train staff, and measure the human risk surface over time. People are the most exploited attack surface, yet awareness training is too often generic box-ticking, disconnected from the real threats an organisation faces. Athena turns your people from the weakest link into a prepared first line — realistic practice, measurable improvement, on EU infrastructure. It is a vision for now, not an available product.

Automated pentesting — working name TBD (roadmap)

This product will safely and automatically probe your own systems the way an attacker would: continuous, guard-railed penetration testing that finds weaknesses before adversaries do. Manual pentests are expensive, point-in-time and scarce; most organisations are tested far too rarely for how fast their systems change. Attacker's-eye findings, hard safety guardrails, a human always in control. The name is still open — my pick from the pantheon is Odysseus, the master of cunning who finds the way in. We will confirm before anything ships.

RunBiz — EU-first CRM to ERP (concept)

RunBiz is a business platform where email is the source of truth. AI reads email threads and auto-creates contacts, companies and deals, and proposes next actions that a human confirms with one click. The CRM was always a data-entry chore that salespeople avoid; the real record of the relationship lives in email and never makes it into the system. RunBiz fixes that — EU data residency end to end, EU-hosted models only (served by Hermes, no US inference), append-only hash-chained audit logs, a per-user view invisible to admins, native Belgian and EU formats. AI proposes; humans decide; nothing is autonomous. It is at concept stage; we use it to tell the vision and validate demand, not to claim features.

In one line: RunBiz turns your inbox into your CRM — with EU-hosted AI and a human always in control.

Project Camiel — the one that isn't about any of this

One project in our family has nothing to do with sovereignty or compliance — and that is exactly the point. Project Camiel is a free, open-source educational game for young children, built to teach the very first building blocks of language and early learning. It is named for, and dedicated to, my father, whose lifelong dream was to make a genuinely good educational game that any child could play, for free.

It is built in Godot, released under an open licence, accessible (WCAG 2.1 AA) from the start, and shaped around gentle, child-friendly play. It is early and evolving — a heartfelt work in progress, not a finished product. We will never sell it or bundle it with the commercial stack. But it shows, in the most human way I know, what “Ethical Innovation” and “European engineering” actually mean to the people behind Euraika: free, open, accessible, made with care. Because it involves children, we hold an extra-high bar on tone and child safety. Always.

How we sound, and what we will not do

Our voice is a differentiator, and I hold the line on it because it is what makes us not “just another AI vendor.” We sound principled, precise, composed, constructive, and European.

We are factual over reassuring. We use active voice, one idea per sentence, calm and measured. We explain how something works, not just that it works. We make responsibility explicit — our software assists, the organisation decides. And we use plain, human language: if a sentence needs a glossary, we rewrite it. The product is simple; the copy must be too.

What we will not do: hype, hyperbole, or vague superlatives. Fear-mongering or regulatory FUD. Black-box, “just trust us” framing. Implying that AI replaces human judgement or accountability. Greenwashing or sovereignty-washing — every sovereignty claim must be literally true of the product. And no political or ideological framing; we are neutral, inclusive and respectful, always.

The test for anything we publish: could a precise, sceptical European CISO read this and find nothing to roll their eyes at? If yes, it is us.

What I ask of everyone who builds with us

Trace everything back to one sentence: Euraika builds a sovereign AI stack for Europe — powerful enough for the most regulated industries, simple enough to use without a manual. Two promises, held together — sovereign and effortless. Neither is optional. If an idea cannot trace back to that, it is off-strategy.

Keep us honest, because honesty is literally one of our values. Treat “top,” “best” and “conquer” as the ambition they are, not as present-day fact. Frame roadmap products as vision, not availability. Protect the names and personal stories of the people here until they have given their consent. Say “supports compliance,” never “makes you compliant” — compliance belongs to the customer, evidenced by our software. And confirm every hard number before it is published.

That is the whole of it. We chose to do this properly. We respect the people who use our software and the continent we build it in. We earn trust with evidence, and we make the responsible choice the easy one. Proof, not promises.

— Bert Colemont

Founder & Visioneer, Euraika