sopra steria next

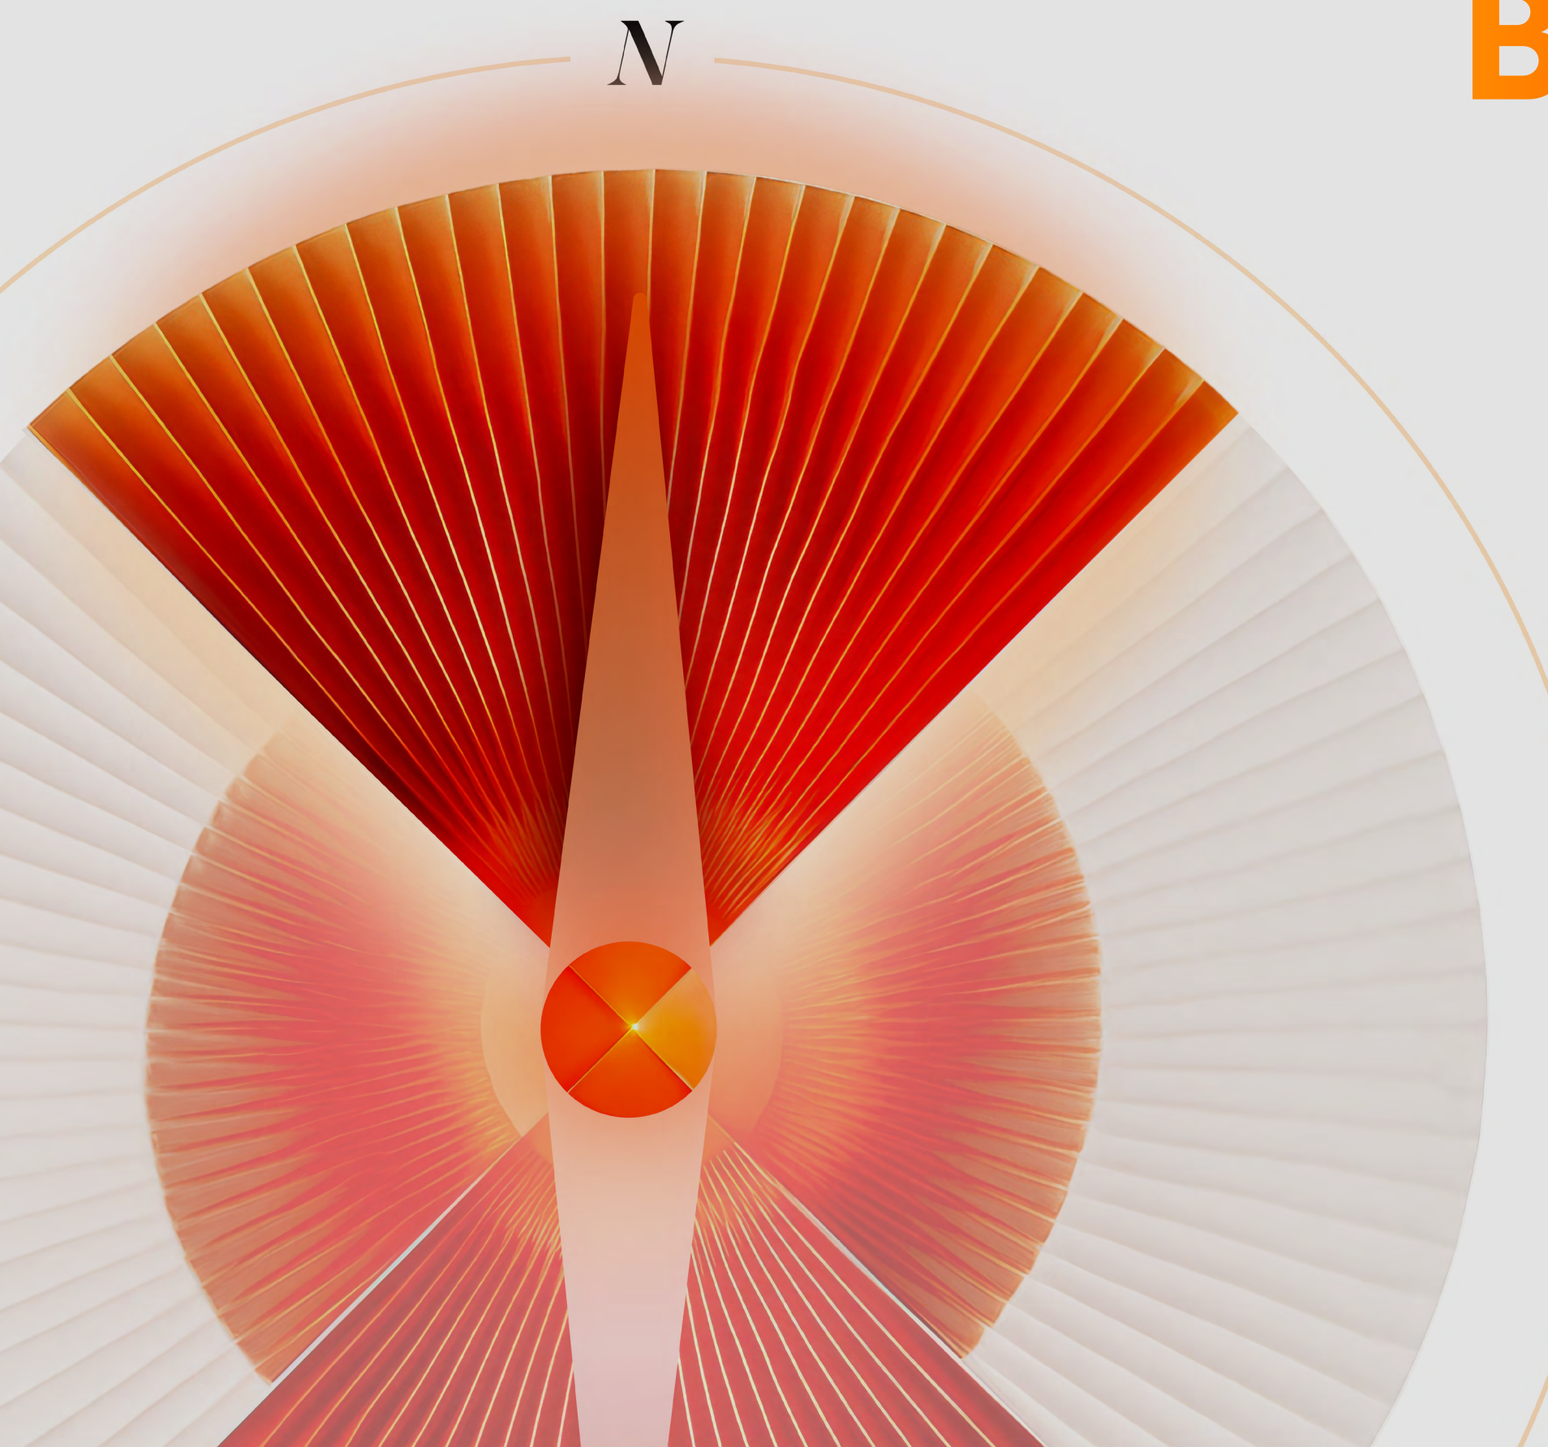

Moving Beyond the GenAI Hype

Industrialising AI with Control, Purpose and Scale

CIO Compass Series

Introduction

The End of *Experimentation*

Artificial intelligence has entered a decisive new phase. The period of curiosity, isolated demos, and fragmented proofs of concept is firmly behind us. Generative AI has graduated from laboratory experimentation to board-level expectation, shifting the core question from *"What can GenAI do?"* to *"How do we turn AI into stable, secure, and repeatable value?"*

This shift reflects a deeper truth: AI is not a mere technical plug-in or a software upgrade. It represents a structural change in how organisations design processes, operate information systems, manage risks, and create economic value. Yet, a significant gap remains between ambition and execution and only a minority of use cases have successfully bridged the chasm to full industrialisation.

We need to be explicit about why AI is introduced, which problems it will solve, and what outcomes it must deliver. We are now entering the **Era of AI-by-Design**, where everything starts with purpose, but this ambition cannot be achieved without a solid foundation. In this new paradigm, success is no longer about the speed of adoption, but the strength of the foundation. It requires **robust governance** that goes beyond simple compliance to align actively with ethical and security requirements. It demands **high-quality, well-governed data** strategies capable of taming vast unstructured information. Crucially, it calls for **flexible, sovereign architectures** that prevent single-vendor dependency, alongside strict **economic discipline** to manage unpredictable inference costs. Ultimately, it relies on **cultural readiness,** moving the workforce beyond basic literacy toward deep adoption and new role definitions.

The organisation capable of mastering these interconnected dimensions will be the ones that convert AI potential into measurable, sustainable performance.

AI Beyond Generation – A new enterprise operating model

Generative AI is only the most visible expression of a broader transformation. True enterprise AI maturity does not emerge from a single technology, but from **four complementary pillars**[1] that redefine the operating model.

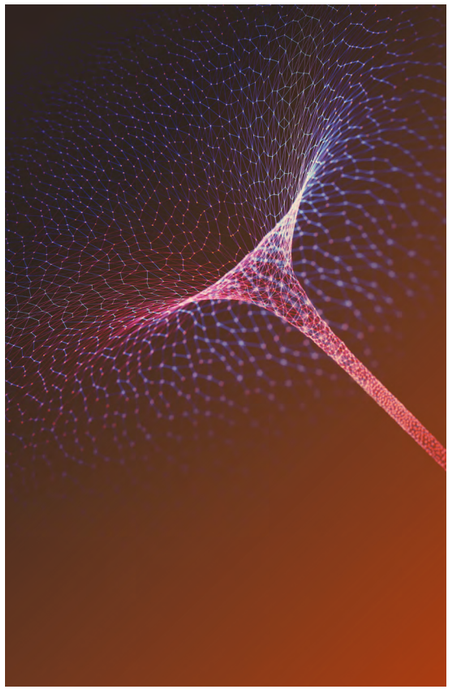

AI for Humans
– Augmented Work
Copilots and assistive intelligence act as force multipliers, enhancing decision-making, accelerating expertise and shifting employees from execution to supervision and validation.

AI for Process
– Agentic Workflows
AI shifts from simple automation to agents that reason, plan, and execute multi-step tasks autonomously, making operations more resilient, adaptive, and predictive. This is a radical change—and a hard one—because it forces a full redesign of processes, controls, and roles, with humans moving from doing the work to orchestrating and governing autonomous agents.

AI for Machines
– Intelligent Physical Operations
AI bridges the digital and physical worlds, bringing intelligence to robotics, OT and edge devices. It enables safer, more reliable and more energy-efficient industrial environments.

AI for Software
– Accelerated Engineering
AI revolutionizes the SDLC by accelerating specification, testing, documentation, refactoring and legacy modernisation, improving velocity and quality across the entire software lifecycle.

When these four pillars converge, AI evolves from a set of use cases into **an enterprise-wide nervous system** that reshapes value chains, shortens decision loops and embeds intelligence throughout operations.

Smart Lean: Data and AI-powered performance

Our philosophy is simple: move beyond traditional lean or pure AI solution deployment by steering every initiative through value, reinforcing business/IT co-creation, and reinventing processes by combining frugal engineering and industrialization of several smart automation solutions.

Over a 24-month cycle, identify, prove, scale, reinvest, we unlock compounding gains by combining data, AI, and the full digital spectrum from IoT to optimisation platforms.

Deployed for a global retail leader, it enabled to redesign 200 automated processes by waves over 24 months and realised 12% savings.

For a major European energy group, a smart lean program led to cut the downtime for power station maintenance processes by half, therefore saving significant maintenance cost and even better freeing electricity production capacity at no cost.

1 AI Study - 4 business-centric categories to leverage performance

The Scaling Challenge – Why industrialisation is harder than expected

Reaching industrial scale is now the true frontier, yet the obstacles are structural rather than purely technical. Organisations face a **"Cost of Cognition"** challenge, where volatility in token consumption and GPU usage can undermine predictable operating models. Furthermore, fragmented architectures and the "data trap" of unstructured information often stall deployment, while the rise of Shadow AI introduces uncontrolled risks regarding data exposure and intellectual property.

In this context, focusing on isolated "plug-and-prompt" experiments is no longer viable. Organisations must shift toward AI-ready operating models grounded in responsibility, sovereignty, and a clear distinction between what should be bought (embedded SaaS AI) and what must be built (custom competitive advantages).

The five challenges CIOs face most often are:

01

The Buy vs Build dilemma

Enterprises must arbitrate between:

- *embedded AI* in SaaS tools (fast value but little control),
- custom AI (high differentiation but complex to maintain).

The right choice depends on sovereignty requirements, competitive differentiation, data advantage and long-term economics.

02

Vendor lock-in and single-model dependency

Relying on a single LLM provider creates exposure to price increases, declining competitiveness, API instability, and geopolitical or regulatory constraints. Enterprises need **model-agnostic** architectures to maintain flexibility and negotiating power.

03

Weak governance and untraceable outputs

Without a controlled ecosystem, risks are that AI systems:

- produce inconsistent or ungrounded outputs,
- cannot be audited or monitored reliably,
- accumulate risks in sensitive domains (legal, HR, finance, cyber).

Trust collapses rapidly in the absence of structured guardrails.

04

Cost volatility and lack of AI FinOps

Without cost governance, inference usage, token drift and GPU bursts quickly overwhelm budgets.

Costs can change month to month and vary dramatically based on model choice and architecture.

05

Shadow AI and new security threats

Scaling AI safely requires security measures tailored to machine reasoning rather than classical IT threats.

Scaling AI safely requires security measures tailored to machine reasoning rather than classical IT threats.

The GenAI Paradox: Massive Adoption, Minimal Production

Moreover, a paradox emerges: while ChatGPT's growth to 700 million users and 18 billion weekly messages confirms its massive adoption, non-professional uses now dominate. The tool designed to transform productivity and business value is increasingly serving as a personal cognitive companion — used more for guidance, information, and writing assistance than for professional automation or innovation[2].

At the same time, fewer than one-third of AI projects ever reach stable production, exposing weaknesses in governance, architecture, and talent[3].

The market's fragmentation - **over 500 GenAI startups launched in 2023** - has created noise rather than clarity. Security and sovereignty concerns are rising sharply, and the shortage of skilled professionals remains acute.

Beyond prompt engineering, organisations must master new disciplines: **AI governance, MLOps, data ethics, FinOps, and change management.**

For CIOs, the mandate is clear: move beyond "plug-and-prompt" experimentation and **build AI on foundations of trust, traceability, transparency, and ROI.**

How people (really) use ChatGPT

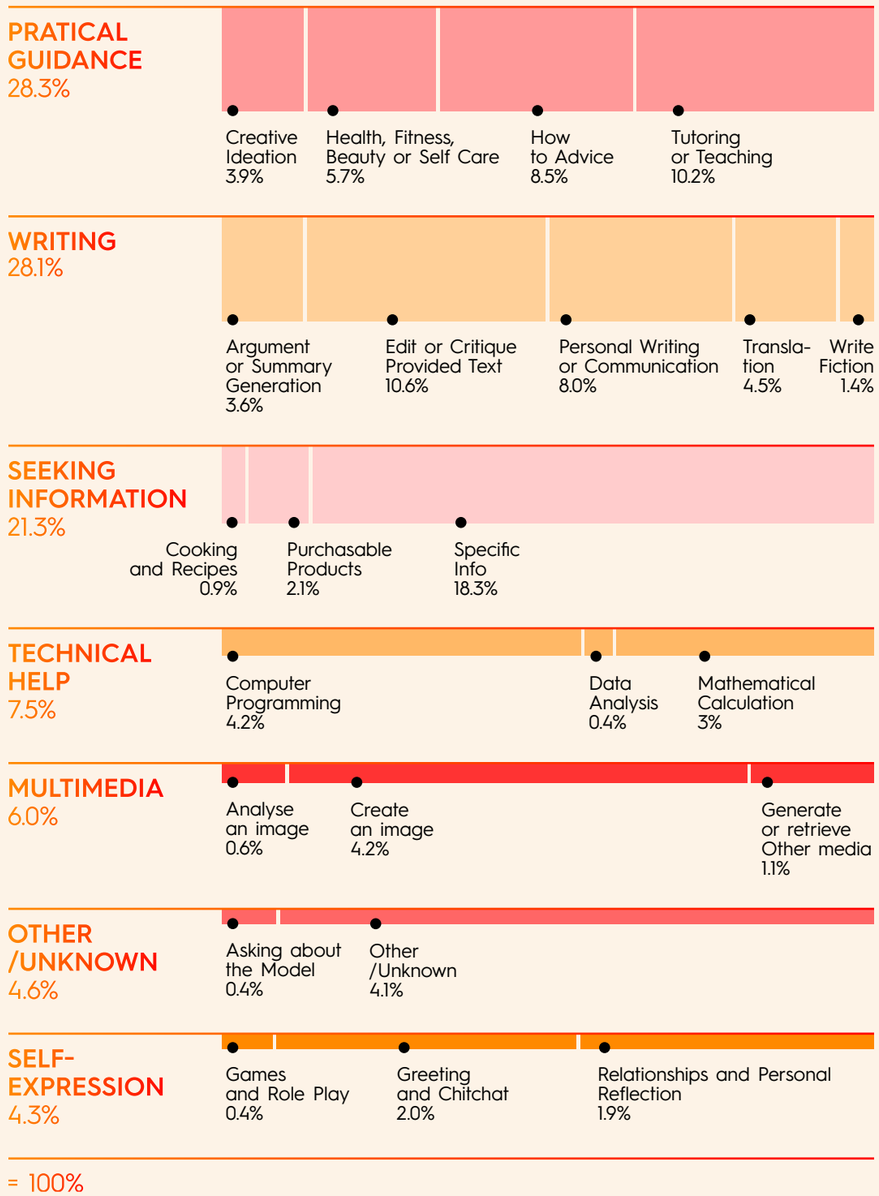

2 Non-professional uses clearly dominate, accounting for 73% of conversations compared with just 53% a year earlier. Three main themes capture nearly 78% of all exchanges: practical guidance, information seeking, and writing support. *How people are using ChatGPT | OpenAI*

3 Gartner Predicts 30% of Generative AI Projects Will Be Abandoned After Proof of Concept By End of 2025

Transform, industrialise, perform: the promise of AI Factory

From banking to aerospace and energy, leaders such as BNP Paribas, Airbus and EDF have adopted the AI Factory model to scale and industrialise AI.

Since 2017, Sopra Steria Next has been supporting a French public service in a transformation journey that unfolded in two key phases. The first phase focused on structuring and industrialising the client's AI platform while opening it up to a wide range of business areas. This involved training subject-matter experts—who have since become AI product managers—and building multidisciplinary development teams, effectively breaking down silos between data scientists and developers. This approach accelerated delivery and ensured that AI solutions could be adapted to specific operational needs.

Once these foundations were firmly established, the second phase aimed to generate large-scale impact, backed by strong commitment from General Management. This dynamic has notably materialised through the deployment of an initial AI tool that enables citizens, assisted by conversational AI, to verbally describe their profile and expectations in order to receive tailored recommendations. Today, with AI positioned as a true strategic lever, use cases are prioritised based on measurable value and validated through lean experimentation.

The next frontier is autonomy and culture: expanding self-service, strengthening business ownership of data, and promoting sovereign, accessible technologies. The approach is iterative and value-driven — embedding efficient, inclusive AI across the organisation.

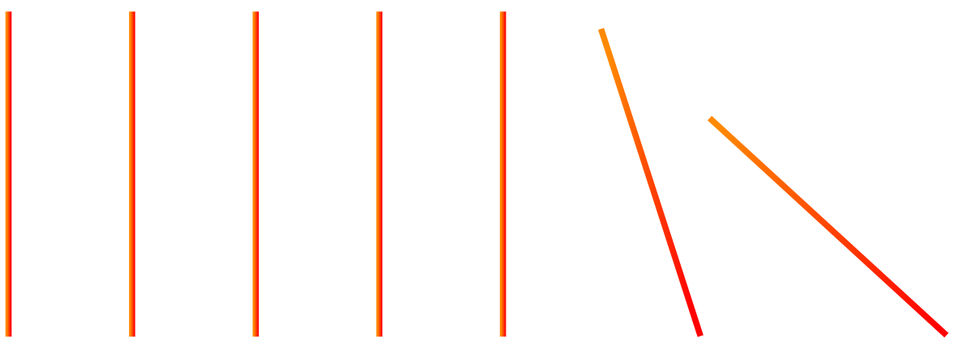

Foundations of Scalable AI The Enterprise AI Blueprint

Sustainable AI requires a cohesive architecture built on three reinforced pillars:

Value & Governance – From Compliance to Competitive Edge

Governance must evolve from a defensive necessity to a value driver. This means establishing risk-tiered gateways that fast-track low-risk innovations while heavily auditing critical systems. Crucially, the strategic focus must broaden from bottom-line efficiency to **top-line revenue generation,** using AI to create new premium services and hyper-personalised experiences. This pillar also encompasses active defense against adversarial threats, such as prompt injection and model theft, ensuring trust remains unbroken.

Technology & Strategy – The "Composite" and "Agnostic" Estate GenAI is only as good as the data it learns from and the architecture that hosts it. Forward-thinking enterprises are moving toward **Model Agnosticism,** utilizing architectures that allow them to swap underlying models as performance and prices fluctuate, thereby avoiding vendor lock-in. This strategy also embraces **Frugality by Design,** deploying highly specialised Small Language Models (SLMs) for specific tasks to reduce latency, carbon footprint, and costs, rather than relying solely on massive, energy-hungry models.

Culture & Skills – Making AI Everyone's Capability

As AI permeates every process, literacy becomes universal. However, this shift requires more than training; it demands role reinvention. We are seeing the rise of "AI Product Owners" and "Process Architects" who bridge the gap between technical capability and business outcome. Success scales not when AI is powerful, but when it is understood, trusted, and seamlessly embedded in day-to-day routines.

From Use Cases *to Process Reinvention*

Most early GenAI deployments focused on automating discrete tasks. The next wave is about transforming the process itself. This involves a shift from reactive operations to predictive, self-adapting systems.

By creating **Composite AI** workflows, blending the creativity of GenAI with the reliability of rule-based systems, companies can collapse cycle times, augment human judgement, and reshape value chains end-to-end. This is where AI meets the strategic agenda: not just doing things faster, but doing entirely new things.

AI as an Industrial Discipline *– Avoiding the Pilot Graveyard*

As the AI landscape matures, leading organisations share a common realisation: the real challenge is no longer discovering what AI can do, but establishing the conditions under which it can operate reliably, sustainably and at scale. This requires treating AI not as a series of exploratory initiatives, but as an **industrial discipline** that must be governed, measured and continuously improved.

To move from scattered pilots to enterprise-wide deployment, organisations must adopt a disciplined and structured model anchored in five principles:

01

Prioritise with conviction

AI initiatives should be tied to explicit value convictions, whether in cost reduction, revenue generation or risk mitigation. This shift prevents experimentation for experimentation's sake and ensures that resources are directed toward opportunities with measurable impact.

02

Build AI Factories rather than isolated experiments

AI Factories bring together standardised tooling, model governance, data pipelines, MLOps capabilities, PromptOps practices, and a secure operating environment. This enables repeatability, reduces technical debt and ensures that new use cases can be delivered faster and at consistent quality.

03

Apply frugality by design

In an era where model costs and GPU consumption can grow exponentially, organisations must actively choose the smallest effective model, optimise inference paths, and benchmark alternatives. Frugality is both a financial discipline and a sustainability imperative.

04

Embed guardrails as code

Security, compliance, data minimisation and lineage verification must be automated within pipelines, rather than relying on manual reviews. This is essential for scaling AI safely and ensuring that outputs remain consistent, auditable and aligned with regulatory expectations.

05

Lead change as a strategic agenda

Adoption does not occur organically. It requires training, new role definitions, communication strategies, incentives and governance alignment across business, IT, security, HR and legal. AI transformation is as much organisational as it is technological.

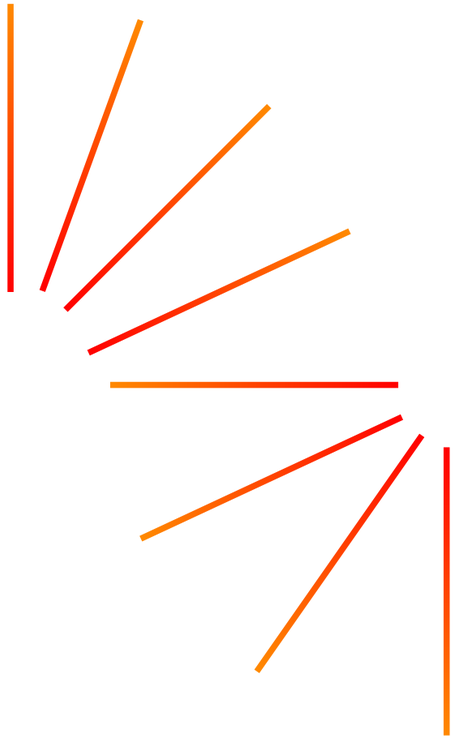

Organisations that internalise these principles will avoid the “pilot graveyard” and transition toward repeatable, scalable, secure and economically viable AI adoption. Industrialisation is not a race for speed but a test of discipline.

Conclusion

From Hype *to Performance*

The transition from GenAI hype to industrial performance is not a technology challenge; it is a **CEO mandate for structural change.** The leaders of the AI-augmented enterprise will not be those who rush the fastest, but those who build with the greatest discipline.

This discipline hinges on shifting the organizational mindset from pilot exploration to **Systemic Intelligence.** The future belongs to those who successfully move from building isolated AI use cases to architecting an **Intelligent Nervous System** where data, governance, and business outcomes are seamlessly unified. The true promise of AI is not automation; it is **augmentation:** elevating human capability and redefining performance thresholds. By combining the creativity of generative models with the rigor of composite, agnostic architectures, organisations can finally turn strategic intent into sustained operational reality, securely and at scale.

Executive Boxes

AI Myths CIOs must now resist

- **Bigger models do not guarantee better performance**
  - Most enterprise tasks are better served by smaller, faster, cheaper models.
- **GenAI will not replace traditional AI**
  - Performance comes from hybrid intelligence, not single-model paradigms.
- **Scaling AI is not primarily a technical problem**
  - The bottleneck is governance, data and organisational readiness.
- **SaaS vendors cannot "solve AI" for the enterprise**
  - They accelerate adoption but increase lock-in and reduce sovereignty.
- **A successful POC does not guarantee industrialisation**
  - Most POCs fail without a dedicated operating model and cost discipline.

What leaders are doing now – The moves that create impact

- **Building model-agnostic AI platforms,** able to switch between models seamlessly.
- **Moving from POCs to AI Factories,** with standardised tools, governance and deployment pipelines.
- **Treating AI as a profit engine,** embedding intelligence in products and processes.
- **Applying frugality by design,** using SLMs and cost-per-inference controls.
- **Enabling controlled AI sandboxes,** fostering safe experimentation for business teams.

Hidden Traps to Avoid

- **Invisible lock-in,** created by over-reliance on a single model provider.
- **Cost drift,** caused by uncontrolled token usage and GPU bursts.
- **Ungrounded outputs,** which persist without retrieval or enterprise context.
- **Shadow AI,** exposing confidential data to public models.
- **Skills illusion,** assuming adoption comes naturally without structured training.

Contacts

Fabrice ASVAZADOURIAN
CEO of Sopra Steria Next
→ fabrice.asvazadourian@soprasterianext.com

Mohammed SIJELMASSI
Chief Technology Officer
→ mohammed.sijelmassi@soprasteria.com

Contributors of our European Network

Olivier BROUSSEAU
Architecture & Tech
→ olivier.brousseau@soprasterianext.com

Isabelle PONS
Data & AI
→ isabelle.pons@soprasterianext.com

Valérie THIBLET
Data & AI
→ valerie.thiblet@soprasterianext.com

Axelle ROESCH
CIO Advisory
→ axelle.roesch@soprasterianext.com

Thomas OTTO
CIO Advisory
→ thomas.otto@soprasteria.com

Philipp MAIER
CIO Advisory
→ philipp.maier@soprasteria.com

Simon KAYE
Architecture & Tech
→ simon.kaye@soprasteria.com

Ray BAKER
Architecture & Tech
→ ray.baker@soprasteria.com

Andy WHITEHURST
Architecture & Tech
→ andy.whitehurst@soprasteria.com

Eirik HASLESTAD
CIO Advisory
→ eirik.haslestad@soprasteria.com

Jahn-Fredrik SJOVIK
CIO Advisory
→ jahn-fredrik.sjovik@soprasteria.com

Tor NESET
Architecture & Tech
→ tor.neset@soprasteria.com

Marius SANDBU
Data & AI
→ marius.sandbu@soprasteria.com

Robert HEXSPOOR
Architecture & Tech
→ robert.hexspoor@soprasteria.com

Hugo HERMSEN
CIO Advisory
→ hugo.hermsen@soprasteria.com

Juan Antonio MARTINEZ SANCHEZ
Architecture & Tech
→ juanantonio.martinezsanchez@soprasterianext.com

Mauro PALMARINI
CIO Advisory
→ mauro.palmarini@soprasterianext.com

Conception: Spintank, photos credits: GettyImages ©