Lumen
Enlight Ideas

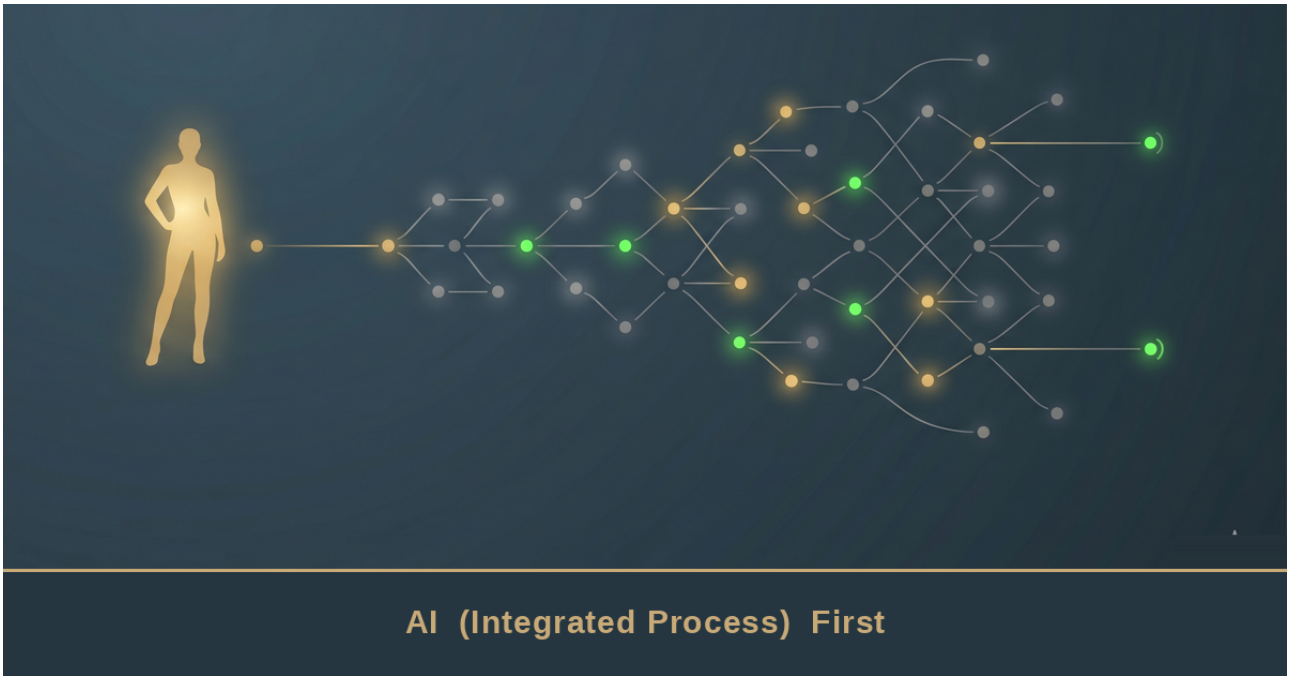
STRATEGIC WHITEPAPER

AI (Integrated Process) First

Why AI-first is a structural and human decision — not a technology one.

Lumen Holding SA T3 • June 2026

The number that frames the moment



The human is the origin and the compass of an autonomous, AI-integrated process.

95% of enterprise AI pilots returned nothing last year — not because the technology failed, but because the structure around it did. "AI-first" is too blunt to be useful: the operating reality is **AI-Integrated-Process-First** — rebuilding the process around the capability, then rebuilding the people around the process. The models are the fulcrum, available to everyone; experience and culture are the lever that moves markets. The advantage is **structural and human**, not technological.

95%

Of enterprise AI pilots returned nothing (MIT NANDA, 2025)

$30–40B

Enterprise spend behind that result

Fulcrum vs lever

Models are the fulcrum; culture & experience are the lever

In 2025, MIT's NANDA initiative reported that 95% of enterprise generative-AI pilots delivered no measurable return, despite $30–40 billion in spend. Only about 5% achieved rapid revenue acceleration. The reflex is to read this as an AI problem. It is not. MIT's own conclusion is that the barrier is integration and learning, not infrastructure, regulation, or talent — and the single factor most correlated with returns was workflow redesign.

The detail matters because it tells you where the 5% won. They did not buy bigger models or larger licences. They bought capability from specialists rather than over-building internally, they redesigned the workflow rather than decorating it, and they wrapped governance and feedback loops around the

system so it improved with use. Their advantage was not a better tool. It was a better operating decision.

The technology worked. The structure around it did not. That distinction is the whole argument of this paper — and it is the reason most of the spend evaporated while a minority compounded.

"AI-first" tells you nothing about what to change. The honest, longer version carries the whole thesis: AI-Integrated-Process-First. Adding a copilot to an unchanged workflow is tooling. Redesigning the workflow, the decision rights, and the economics around the capability is structure. Only one of the two creates value, and the 95% failure rate is the evidence — companies bought capability and pointed it at an unchanged process run by unchanged people.

AI-first is a slogan. AI-Integrated-Process-First is the operating model

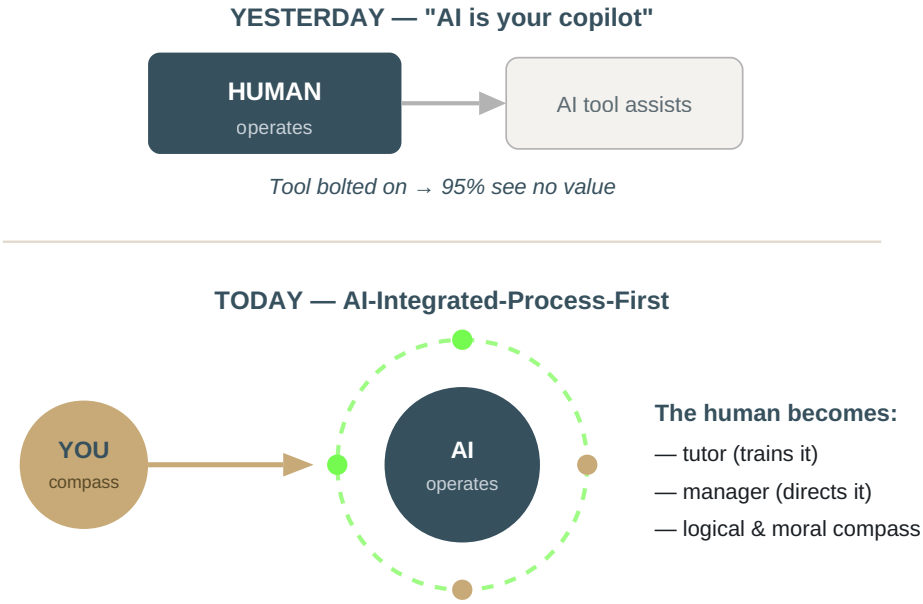
The advantage does not come from adding AI to a process. It comes from rebuilding the process around AI — and then rebuilding the people around the process.

A concrete example makes the difference obvious. Bolt an AI assistant onto a claims process that still routes every case through the same four approvals, and you get a faster first draft and the same cycle time — the bottleneck never moved. Rebuild the process so the model triages, drafts, and clears the routine cases while a human owns the exceptions and the policy, and the cycle time collapses. Same model, same data. The difference is entirely in the structure around it.

More than half of enterprise GenAI budgets went to sales and marketing tools, while the largest measurable returns sat in back-office process redesign. Meanwhile a shadow AI economy has formed: roughly 90% of employees already use personal AI tools daily, while only about 40% of companies have official deployments. The capability is already inside the building. What is missing is the structure to convert it into value.

The inversion

Who flies the process — and who is the copilot



From "AI assists the human" to "AI operates; the human is the tutor, manager and compass."

The destination is a hybrid, human–agentic company

The endpoint is not "humans using AI." It is a hybrid organization — part human, part agentic — whose internal dynamics start to resemble those between people: delegation, trust, supervision, escalation. You will not manage agents like software; you will manage them like a team. That means the work is not installing tools — it is redesigning processes and people's skills until the two species work together. Processes and tools must re-adapt to that reality, not the other way around.

The inversion: who is the copilot now

LUMEN CEO SAYS

For two years, AI was sold as our copilot. In a process built around AI it is the opposite: the AI operates — and we become its copilot, its tutor, and its logical and moral compass. That inversion, not the tools, is the real work.

Riccardo Rolando — CEO, Lumen Holding SA

Reskilling here does not mean teaching everyone to write prompts. It means moving people up the value chain: from doing the repetitive work to specifying it, supervising it, and owning the judgement the machine cannot. The scarce skills become problem framing, exception handling, quality control, and the moral and commercial judgement about what the system should and should not do. Those are human skills, and they have to be deliberately built — they do not appear because a licence was purchased.

This is why becoming AI-first is a human decision before it is a technical one. The organizations in the successful 5% redesign workflows rather than decorate them, buy capability from specialists rather

than over-build internally, and wrap governance and feedback loops around the system so it improves over time. None of that is a technical achievement. It is an operating-discipline achievement.

The services industry is being repriced

Nowhere is the shift clearer than in IT services, consulting, and system integration — a model historically priced on headcount and billable hours. Demand is now explicitly conditional on AI: IBM's 2025 research found 86% of consulting buyers actively seek AI-enabled services, and 66% would stop working with firms that fail to incorporate them. The incumbents are restructuring in real time — Accenture cut roughly 11,000 roles while committing $3 billion to AI and pledging 80,000 AI-focused hires.

The mechanism is simple and brutal. The classic services model is a pyramid: leverage many junior hours under a few senior partners and bill the difference. When AI compresses the junior hours, the pyramid loses its base — and with it the margin. A firm that keeps selling hours will watch its price per outcome fall. A firm that sells a governed, AI-integrated outcome — priced on the result, not the effort — keeps the value it creates.

When the work that justified large teams can be compressed, value migrates from effort to structure: from how many people you can deploy to how well you can design, govern, and operate an outcome. Firms that sell hours will be repriced. Firms that sell governed, AI-integrated outcomes will not.

Cybersecurity: the gap turns adversarial

> ***Autonomous attackers reach domain dominance in under an hour. A periodic, manual red team cannot keep pace.***

If AI-first is a structural question, cybersecurity is where the cost of getting it wrong becomes immediate. An MIT study showed an AI model reaching domain dominance on a corporate network in under an hour, with no human intervention, adapting tactics on the fly to evade endpoint detection. Autonomous attack frameworks are becoming a defining capability of cybercriminals, and a state-backed group has already automated attacks against roughly 30 organizations.

The economics are the real shock. Automation collapses the cost of mounting an expert-grade attack, so the volume and sophistication available to a low-resource actor rise sharply. Defence faces the same automation — but only if it is structured to use it. An organization that runs a once-a-year manual penetration test is bringing an annual cadence to a continuous fight, and only about 35% of organizations use advanced machine learning for threat hunting at all.

The traditional red team — periodic, manual, expensive — cannot match the cadence of automated offense. The right structure is not "AI replaces the red team" but AI amplifies it: agents perform the broad, repetitive enumeration; humans drive the creative, high-value exploitation and the judgement about business impact. Defenders who treated security as a tool purchase lose; the ones who treat it as a continuously tested structure keep up.

Fulcrum and lever

The models are extraordinary — and available to everyone. The big players, the model makers and hyperscalers, are the fulcrum. A fulcrum moves nothing on its own. The lever is what you place on it: institutional experience, market context, and culture. That is the part that cannot be downloaded, and it is the part that lifts markets.

This is why the heaviest spenders are over-represented among the failures. Buying more of the fulcrum — more models, more seats, more pilots — does not lengthen the lever. It produces motion without leverage: activity, dashboards, and announcements with no movement in the P&L. The firms that win spend less on the fulcrum and far more on the lever: on redesigning how they work and on the people who will run the hybrid system.

The fulcrum is available to everyone. The lever — experience and culture — is not.

What Lumen is building

This is the transition Lumen is built to lead. We are assembling an Adoption & Change Management engine — technology used fluidly, but inside a framework carried by subject-matter experts — so an organization can become AI-Integrated-Process-First without losing what makes it valuable. In practice that means redesigning the high-value processes first, defining where agents operate and where humans hold judgement, building the governance and feedback loops, and reskilling the people who will run the result. The dream behind it is specific: let people train their own agentic company while keeping the know-how and market context that are theirs alone.

Over the next three to five years, advantage will not accrue to whoever owns the most models or the largest tool budget. It will accrue to whoever can integrate capability into a governed operating system — workflows redesigned, decision rights realigned, security treated as a continuously tested structure, economics priced on outcomes. AI-first was never a technology decision. It was a structural and human one. We don't follow innovation. We structure it — and we bring the people with us.

Key takeaways

- → AI-first is structural and human, not a technology purchase.
- → Rebuild the process around AI, then rebuild people's skills around the process.
- → AI operates; the human is its tutor, manager, and logical & moral compass.
- → Models are the fulcrum (available to all); experience and culture are the lever.

Lumen
Enlight Ideas

Lumen Holding — an international platform headquartered in Switzerland, operating at the **intersection** of technology, capital and real assets. We don't follow innovation; we **structure** it. www.lumen.holdings