The AI-empowered organization

A human-centric and agentic transformation approach for closing the gap between AI ambition and frontline execution

eraneos

Contents

Executive summary

The AI investment cycle has hit a paradox.

Boards have allocated budget, executives report high confidence, and most large organizations now run dozens of AI initiatives, yet 95% of those pilots produce no measurable return, and only 5% reach production. The standard explanation is technical: the data isn't ready, the architecture is wrong, the use cases were poorly chosen.

Our research suggests something different. Across European organizations we surveyed for the Eraneos People & AI Study 2026, 68% of C-level executives express confidence in AI taking real action on behalf of their organization. Among senior specialists, that figure drops to 21%. That gap of almost 50 points is not an artifact of communication. It is the structural reason most AI programs stall: the people closest to the work, and best placed to redesign it around AI, do not yet trust the systems they are being asked to adopt. The study measures the full shape and cost of this problem; this paper sets out the AI-empowered organization that closes it.

This has a consequence most consulting frameworks miss. Trust in AI does not flow top-down. It is built bottom-up, through frontline enablement, calibrated authority, and shared accountability between humans and agents. Programs that invert the standard playbook, investing first in the practitioners who will work alongside AI rather than in central control structures designed without them, convert executive ambition into frontline adoption measurably faster. In this model, People & Enablement is not a downstream training workstream. It is one of two structural bridges between the value and trust engines – together with Strategy & Value – and the one that determines whether the rest of the operating model actually lands.

The AI-empowered organization is the operating model that makes this work. It rests on two engines: a value engine that turns AI ambition into measurable enterprise outcomes, and a trust engine that engages directly with the risks that matter most in the agentic era: error compounding across agent chains, tool-use exposure, prompt injection at scale, and accountability when an agent acts on a human's behalf.

Both engines move in parallel through three phases: Augment (AI accelerates existing work), Reinvent (AI transforms how work happens), and Scale (AI enables what was not previously possible).

The case studies in this paper illustrate the pattern: one engagement converted board-level AI ambition into a five-year roadmap with a $50M+ annual savings path and 30 prioritized use cases for 2026 alone; another grew a struggling internal GenAI platform from low adoption to over a thousand active users in six months by inverting the standard top-down playbook. Both follow the model set out in this paper.

This paper sets out the model, the journey, and the design choices that determine success. It draws on our research into how AI capability is built inside organizations, on two client engagements, and on the operating-model work we have done across financial services, insurance, energy, public sector, and industrial clients. It closes with the diagnostic we use to map clients to their starting point – and the destination they should be aiming for.

The prize is bigger than the outcome of any single AI program. The AI frontier moves every few months; capabilities that did not exist eighteen months ago are now operationally central, and the cycle is accelerating. Organizations that treat each wave as a discrete change program burn out the same people on each new program. Organizations that build a sustained AI capacity to absorb continuous change – distributed across the operating model, not centralized in a transformation office – compound advantage over those that do not.
That capacity is what the model in this paper is ultimately designed to produce.

---

[1] MIT, State of AI in Business, 2025

[2] People & AI Study – Levels of AI adoption, capability, trust & governance in European organizations, 2026

The CEO's AI transformation challenge: ambition has outrun execution

Most AI conversations in 2026 are not about whether to invest. The case for AI is settled, budgets have been allocated, and most large organizations now run dozens of initiatives. The harder question is the one CFOs are starting to ask: where is the impact on the bottom line?

The honest answer is that ambition is converting into activity, but activity is not converting into earnings. Pilots ship. Tools get rolled out. Adoption metrics rise. But the line items that the AI investment case was supposed to move (cost-to-serve, cycle time, revenue per relationship, gross margin on operations) are moving slowly or not at all. AI transformation is failing not because organizations lack ambition, but because that ambition is being articulated at a level, and in a language, that the people who would have to make it real do not share.
Our People & AI Study makes this visible: in nearly every organization we have surveyed, executives describe AI as a strategic priority while the people closer to the work describe it as something happening to them, not with them. The investment goes in at the top. The behavior change that would convert it to P&L impact does not happen at the bottom.

This produces three patterns we see consistently across our client base, each with a direct cost.

1. **The pilot trap**
   Organizations spawn dozens of initiatives that prove technical viability but never reach the P&L. The use cases work in demo. The deployment doesn't compound. Productivity gains stay trapped inside individual workflows rather than aggregating into measurable functional cost reduction. The reason is not technical readiness: no one outside the AI team has been equipped to absorb the new way of working into their daily process, and the gains evaporate when the pilot team moves on.

2. **Transformation fatigue eating the run-rate**
   In financial services in particular, four years of overlapping change programs (cloud migration, agile transformation, digital banking, ESG reporting) have left middle management exhausted. AI lands as the latest in a series rather than the moment that changes the trajectory.

The cost shows up as initiative drag: the same change muscles as initiative drag: the same change muscles that delivered the previous waves are being asked, again, to deliver this one, with each additional program producing diminishing returns and longer time-to-value.

3. **Trust-bound sectors stalling on usage**
   In public administration, healthcare, and parts of utilities, regulators are not the bottleneck. Internal trust is. Where employees cannot explain how an AI system reached a decision, they will not endorse using it for decisions that matter, regardless of what the compliance documentation says. The output is generated. It does not make it into the workflow. The cost is silent but real: organizations pay for AI capacity they then route around.

These patterns share a root cause that explains why the bottom-line impact has not arrived. Organizations are deploying AI into operating models that were designed for a world in which humans were the only execution layer. That model cannot absorb agentic AI, and it cannot convert AI capability into earnings without being redesigned. The redesign starts not with the architecture diagram but with the people who will work alongside the agents.

Our research shows: AI is ready. Most organizations are not
The Eraneos People & AI Study 2026 – a European survey across financial services, transport & logistics, energy & utilities, and the public sector – measures the gap between executive and frontline confidence in AI in precise terms, sector by sector, and quantifies what that gap costs organizations every year it persists. Two observations from that research shape the rest of this paper.

The first builds on the figure in the executive summary. The gap between executive and frontline confidence is not anecdotal, and it is not confined to one question. It holds across every dimension we measure: how clear the rules are, how trusted the outputs are, how embedded the capability is, how aligned leadership and the workforce are on what is actually working. Where executives see a program on track, frontline workers consistently describe something different.

The second concerns how that gap closes. When we cut the data by how AI capability is built in an organization, the differential between methods is dramatic. Programs that develop AI capability through occasional training alone report frontline confidence levels in the high teens. Programs that integrate AI into everyday work, with role-specific examples and time allocated for hands-on experimentation, see frontline confidence rise into the low fifties. Passive support barely moves trust; active participation transforms it. Standalone training, the centerpiece of most enablement programs, delivers the weakest returns of any structured approach we measured.

The implication is sharp, and it shapes the rest of this paper. Treating enablement as the last step of an AI program is the most common reason those programs fail to land. Whether AI lands in the P&L depends on closing the distance between what executives expect and what their frontline is equipped to deliver. That distance does not close on a schedule set by the steering committee. It closes when frontline workers gain hands-on time, contextual examples, and AI integrated into the work they already do, in parallel with the governance and capability work that needs to happen alongside.

This is why the operating model we set out in the next section places people and enablement at the structural center, not as a downstream workstream. The full study will go deeper into how the gap takes shape across sectors and what it costs; this paper sets out the operating-model implications.

Two engines for closing the gap: value and trust

Becoming an AI-empowered organization is not about deploying more tools. It is about redesigning how the organization creates value with AI as a core part of the operating model, not an add-on to it.

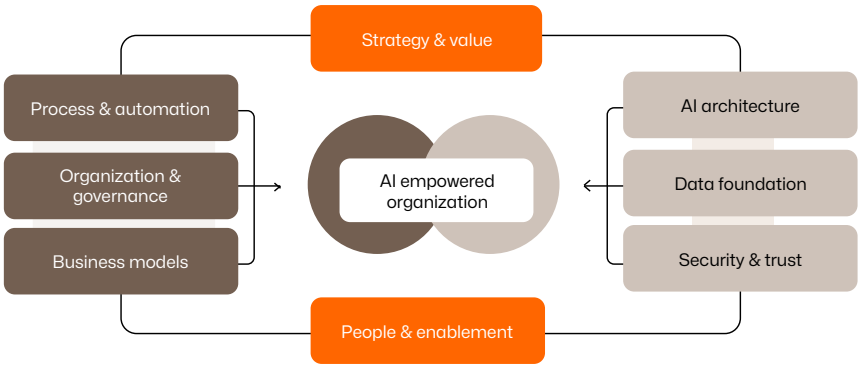

Figure 1: An AI-empowered organization consists of two engines and their eight building blocks

The model we use with clients rests on two engines. The **value engine** translates AI ambition into measurable enterprise outcomes. The **trust engine** allows AI to operate safely and scale sustainably in the agentic era. Each engine has four building blocks. Together they form the complete operating model.

The value engine: from ambition to enterprise outcome

Strategy & value
AI investments compound only when they roll up to a CEO-level agenda tied directly to the P&L. We work with executive teams to define where AI changes how the organization competes, not where it improves a process. The output is a small set of strategic bets with explicit value targets, sequencing, and accountability, replacing the long pilot inventory most organizations carry. In our financial services case below, this approach converted 300+ raw use case ideas into a $50M+ savings roadmap.

Processes & automation
Most organizations have automated tasks. The AI-empowered organization redesigns end-to-end workflows so humans set direction and AI agents handle execution. The defining shift is from human-in-the-loop, where a person approves every step, to human-on-the-loop, where a person supervises outcomes and intervenes when needed. This is not a tooling change; it is a redesign of how authority flows through a process.

Organization & governance
As agents take on execution, the operating model must adapt. Roles, decision rights, and accountability structures need redesign so humans and agents function as a single team. The hardest question (and the one most organizations defer) is who is accountable when an agent acts. We address this explicitly in the trust engine; the answer cannot be left to legal review after the fact.

Business models
The most advanced AI-empowered organizations do not only use AI to reduce cost. They use it to create revenue that was previously out of reach: AI-native products, hyper-personalized client experiences, and proprietary data assets sold as services.

The trust engine: built for the agentic era

The trust engine is where most consulting frameworks become generic. Compliance, auditability, security: the language is familiar, but the substance applies equally to a CRM rollout. The agentic era introduces a different class of risk that requires different responses.

AI architecture
A composable, model-agnostic foundation that can evolve with the frontier matters more in the agentic era than in any previous technology cycle, because the rate of frontier progress is unprecedented. The architecture must also support **observability across agent chains:** when an agent calls another agent that calls a tool, every consulting framework that talks about „auditability" needs to specify what is being audited and across what trace. We design this in from day one.

Data foundation
Every AI agent and every decision is only as good as the data behind it. In the Augment phase, the priority is governed, accessible data for use cases. In Reinvent, data evolves into domain-owned data products, owned by the business teams that use them. In Scale, proprietary data assets become a competitive moat that cannot be replicated by a competitor with the same models. The progression matters: organizations that try to skip directly to „data as a moat" before solving access basics consistently fail.

Security & trust in the agentic era
Four agentic-specific risks define the trust engine's hardest work and must be engineered in from day one to help ensure your organization is resilient by design:

1. **Error compounding** across agent chains, where small reliability failures multiply (a five-step chain at 95% per-step reliability ends at 77%), addressed through bounded contexts and explicit checkpoints.

2. **Tool-use risk,** where an agent connected to real systems can take consequential actions on flawed inputs, addressed through least-privilege permissions, rate limits, and confirmations on irreversible operations.

3. **Prompt injection at agent scale,** where agents reading external content can be hijacked by instructions embedded in it, addressed through segregation of trusted and untrusted inputs and output filtering.

4. **And accountability assignment:** when an agent acts on a human's behalf and produces harm, the question of who is responsible has no satisfying technical answer; it must be resolved as an organizational design choice in advance, not after the fact.

People & enablement
Technology alone does not create an AI-empowered organization. People do. This is one of the two building blocks that bridge the value and trust engines: where Strategy & Value sets the direction both engines deliver, People & Enablement is where both engines meet the daily reality of the work. It converts the value engine's strategic intent into adoption on the ground, and gives the trust engine its human substrate. Five components define the work.

Five components define the People & AI enablement work:

- **AI literacy at the right depth for each role**
  Executives need fluency in what AI can and cannot do strategically. Middle managers need to understand how to redesign workflows around it. Frontline specialists need hands-on capability with the specific tools relevant to their work. Generic enterprise-wide AI training programs underdeliver on all three. Differentiated curricula by role and seniority, with practical application built in, are what convert literacy into capability.

- **Role redesign for human-agent collaboration**
  Most existing job descriptions assume humans are the only execution layer. As agents take on more work, roles change: some shrink, some expand, some emerge. Working through this deliberately, with line managers and HR co-owning the redesign, prevents the quiet erosion of role clarity that derails most AI programs in their second year.

- **Co-design methodology**
  The single most powerful intervention we have measured for closing the executive-to-frontline confidence gap is involving frontline workers in shaping the AI tools they will use. Co-design is not a workshop format; it is a way of structuring the program. The people who will work alongside an agent help define what it should and should not do before any tool is built.

- **Change architecture**
  AI transformation is not change management bolted onto a tech rollout. It is a redesign of how decisions flow through the organization. This requires explicit attention to authority calibration (who can override the agent, when), to escalation paths (what does the agent do when uncertain), and to feedback loops that route practitioner experience back into platform development.

- **AI Ambassador networks**
  Both case studies in this paper rely on the same pattern: a distributed network of practitioners who serve as local enablement nodes inside each business unit. Ambassadors translate central capability into context-specific application, surface practitioner feedback to platform owners, and accelerate adoption through peer-to-peer influence in ways that no centralized rollout team can match. Building this network deliberately, with named roles and protected time, is the structural lever that scales enablement without scaling the central function proportionally.

These five components, executed together, are what makes the bottom-up trust thesis operational rather than aspirational.

The transformation arc: Augment, Reinvent, Scale

Organizations do not become AI-empowered through a single program. They get there through a disciplined journey across all eight building blocks, moving through three connected phases. Each phase represents a meaningful change in how AI operates in the business, and in what becomes economically possible.

We use the Augment / Reinvent / Scale arc deliberately. It maps cleanly to how value compounds: first within roles, then across workflows, then through the operating model itself.

The phases are not gates to be passed sequentially (most clients run all three concurrently in different parts of the business), but each phase has a distinct character, set of investments, and expected return.

Becoming an AI empowered organization is a unique transformation journey

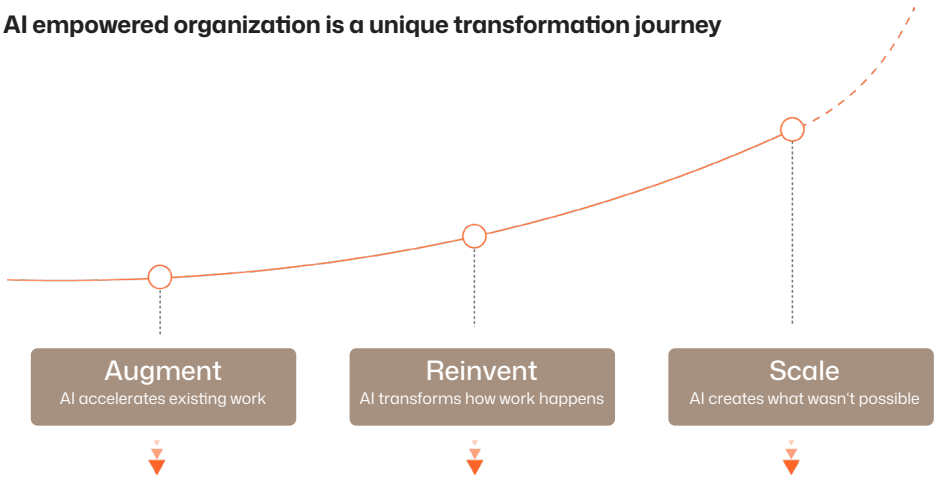

		Augment	Reinvent	Scale
Value Engine	Strategy & value	AI on the CEO agenda, tied to P&L	AI reshapes where & how we compete	AI redefines the industry we play in
	Processes & automation	Every role augmented by a copilot	Agents run end-to-end workflows	Headcount-independent scaling
	Organization & governance	Accountability for AI outcomes is clear	Humans and agents work as one team	Agent-native operating model
	Business models	Margin expansion through efficiency	New AI-powered products and services	Revenue streams previously out of reach
Trust Engine	AI architecture	Trusted platform ships use cases fast	Modular, model-agnostic agent stack	Infrastructure evolves with frontier
	Data foundation	Clean, accessible data for all use cases	Data as a product, owned by domains	A sovereign, non-replicable data moat
	Security & trust	AI is compliant, auditable, and safe	Trust, ethics & resilience built by design	Adversarial resilience at agent scale
	People & enablement	AI fluency reaches every employee	People use AND orchestrate agents	Organization learns faster than market

Figure 2: The transformation arc: three phases plotted across the eight building blocks, with characteristic outcomes per phase.

Phase 1, Augment: AI accelerates existing work

The Augment phase proves value quickly and builds the foundations that make deeper transformation possible. Every role gets a copilot. AI assists with analysis, drafting, research, and routine decisions, compressing the time it takes to complete high-volume, repetitive tasks. The CEO-level AI agenda is established, accountability for AI outcomes is made explicit, and the data and architecture foundations begin to be strengthened.

The outputs are tangible: measurable individual productivity gains in target functions, faster cycle times, early proof points that AI can be trusted, and a growing base of AI-fluent employees. Critically, this is also the phase in which frontline confidence starts to catch up with executive ambition, through active participation rather than training catalogs.

Phase 2, Reinvent: AI transforms how work happens

In the Reinvent phase, the operating model begins to change. Agents run end-to-end workflows: gathering information, coordinating across systems, making recommendations, and executing defined actions, with humans supervising outcomes rather than intervening at every step. This is the shift from human-in-the-loop to human-on-the-loop.

People begin to experience AI agents as colleagues with names, defined roles, and consistent behaviors. New AI-powered products and services emerge. Data moves from a shared resource to a domain-owned asset. The organization's competitive position begins to shift visibly. Material functional cost reductions become achievable, with quality held or improved.

Phase 3, Scale: AI creates what was not previously possible

In the Scale phase, AI no longer accelerates or transforms existing work – it creates what was previously out of reach. Growth becomes headcount-independent: output grows without proportional headcount because agents handle execution at a fraction of the cost of human labor. The operating model becomes agent-native, designed from the ground up for human-agent collaboration.

Data becomes a sovereign, non-replicable competitive moat. Revenue streams that were previously inaccessible become real. The trust engine operates at full capacity: AI is adversarially resilient at agent scale, accountability is built into the architecture, and the organization learns faster than the market.

Two case studies

Case 1: Driving AI adoption from the ground up at an insurance IT services provider (Augment to early Reinvent)

The client is the international IT services arm of a major European insurance group, responsible for the entire IT infrastructure and digital transformation of the parent business. A large share of their technology workforce is engaged in the software development lifecycle. Having recognized GenAI's potential early, they had centralized AI capability by building their own internal GenAI platform.

The initial results fell short. Despite significant investment, adoption among senior developers remained low. The platform focused on generic use cases and showed gaps in context understanding, process traceability, and customization for practical use by product teams. The rollout was further constrained by technical limitations, inconsistent data quality, and limited GenAI expertise across the organization. The pattern matched what we see consistently in experienced-practitioner populations: skilled professionals view centrally imposed AI platforms as something happening to them rather than something they have shaped.

Instead of doubling down on the central platform, the program pivoted to a coaching-first model alongside continued tool development. Four mechanisms drove the shift. An **enablement network** strengthened peer-to-peer learning and hands-on training, building distributed AI expertise where the work happens. **Practitioner-driven improvement** routed real user feedback into platform priorities, sharpening context engineering and enabling agentic workflows via MCP to reach external systems. **Tool-agnostic guidance** built trust by being honest about when to use the internal platform versus specialized third-party tools. And **agentic AI implementation** moved teams from static LLM prompts to self-service multi-agent orchestration for documentation, analysis, and knowledge sharing, with human oversight retained where it mattered.

Within six months, the platform grew from a struggling pilot to over a thousand active users, supported by working sessions, community townhalls, and a growing cohort of AI Ambassadors transferring knowledge across teams. Tens of high-impact features were defined and shipped through a structured development pipeline. Several cost-saving use cases became embedded into daily workflows, and the AI Office itself transitioned from being a centralized platform provider to a facilitator of shared learning – exactly the organizational design shift toward product-led ownership described later in this paper. The bottom-up enablement approach also visibly reduced the risk of shadow AI.

.

The lesson: in a population of skeptical, experienced practitioners, AI adoption did not scale through better central tooling alone. It scaled when the people closest to the work were enabled, listened to, and trusted to drive the transformation themselves, exactly what our research on active participation versus passive training would predict.

Case 2: Translating AI ambition into a $50M roadmap at a European financial services firm (Strategy & Value, with implementation underway)

The client is a leading European financial services company with a clear board-level ambition: significant cost savings through AI, identified in an earlier assessment at tens of millions per year. The ambition was real, but the path to delivery was not. A C-level survey had captured more than 300 AI use case ideas of widely varying maturity and unvalidated business impact. Without operational clarity, the savings risked staying theoretical, another instance of the value gap that defines most enterprise AI programs.

We applied our AI Design Sprint Framework to convert ambition into an executable plan, positioning the work as a strategic transformation across the whole organization rather than a series of departmental experiments. First, an AI maturity assessment grounded the program in fact, mapping more than 30 end-to-end processes to build the analytical base. Second, the 300+ raw use case ideas were translated into a structured backlog, mapped to specific value streams to produce a use case heatmap and identify white spots (areas of the business with high automation potential but low current coverage). Third, AI Design Sprint workshops ran across the prioritized value streams, varying in depth from opportunity mapping to full end-to-end process redesign depending on each stream's potential.

The work accounted for what most roadmaps gloss over. Regulatory constraints in finance and HR required incremental approaches within fixed boundaries. Customer service and marketing offered headroom for broader redesigns where AI could orchestrate across multiple systems. Implementation lead times, dependencies between use cases, and the organizational change curve were all built into the sequencing.

The output was a board-ready five-year AI roadmap with a clearly defined path to over $50 million in efficiency gains: around 30 prioritized use cases targeted for 2026 alone, drawn from a validated portfolio of 430 use cases extending through 2030. The roadmap gave the executive team transparency on when and where each business unit would see change, allowing resources, expectations, and leadership priorities to align before implementation began. The client has since moved from strategic exploration into implementation of multiple AI use cases at scale.

The lesson: AI ambition without disciplined prioritization produces 300 pilots and the 95% failure rate that defines the industry. Disciplined prioritization, grounded in real process data and tied to measurable business impact, is what turns ambition into a value creation plan.

The pattern across Eraneos' core industries

The distance between executive ambition and frontline adoption takes a different shape in each sector we work in.

Financial Services
Regulatory exposure and transformation fatigue make executive conviction the typical bottleneck, not frontline trust.

Transportation & Logistics
Workforce shortages create urgency; the operational consequence of error demands that enablement precede authority expansion.

Energy & Utilities
Safety-critical operations and retiring control-room expertise call for an advisory-to-action progression earned through operator verification.

Life Sciences
Validation, auditability, and reproducibility are structural; the trust engine must be engineered to GxP (good practice guidelines) and submission standards from day one.

Public Sector
Civil servants are accountable to citizens; explainability and algorithmic accountability are political preconditions, making bottom-up trust building essential.

Automotive
A dual-speed AI agenda (in-product AI under ISO 26262 and SOTIF safety standards, enterprise AI driving productivity) must run on one operating model.

Defence
Sovereignty constraints and human-in-the-loop doctrine make the trust engine architecturally non-negotiable; AI must be sovereign, fully auditable, and accountable by design.

How we approach enablement: the Eraneos AI Way-of-Working

The cases above share more than a thesis. They share a methodology, refined across our engagements in Organizational Excellence & Transformation, that turns the bottom-up trust argument into something a program can actually execute.

The method has three operating principles and three signature elements.

Operating principles:

- We **co-design with the people who will work alongside AI** before any tool is rolled out, not as a workshop but as a structural way of running the program.
- We **sequence enablement and authority expansion** deliberately, so frontline workers earn trust through verified experience rather than being granted it by policy alone.
- And we build **governance and operational practice** together, so the guardrails reflect what genuinely works in the daily flow of the work.

Signature elements:

- Every engagement establishes an **AI Ambassador network** of named practitioners across business units, with protected time, defined responsibilities, and direct lines back to the central platform team
- We use **practitioner-driven improvement** cycles that route real user feedback into platform priorities monthly, not annually, sharpening the tools faster than any roadmap document could.
- And we build **role-and-authority maps** that explicitly calibrate where humans intervene, where they supervise, and where agents can act, with the calibration reviewed quarterly as confidence builds.

The method is not a substitute for technical excellence in architecture, data, or security. It is what makes those investments actually compound into transformation rather than stalling at the pilot stage. In the cases above, this is the layer that converted a struggling internal platform into an organization-wide capability, and that turned a $50M ambition into a roadmap the executive team could actually execute.

Designing for shared authority: from centralized control to product-led ownership

If trust travels bottom-up, the organizational form that institutionalizes it matters as much as the methodology that builds it. The AI Way of Working tells us how to develop trust between people and the agents they work with; the organizational design tells us who owns what once that trust exists. Get the methodology right and the structure wrong, and the gains evaporate the moment the program team moves on.

The direction of travel is clear: as AI matures in an organization, responsibility must distribute more widely, moving from centralized control toward domain-led ownership, with appropriate governance at the center.

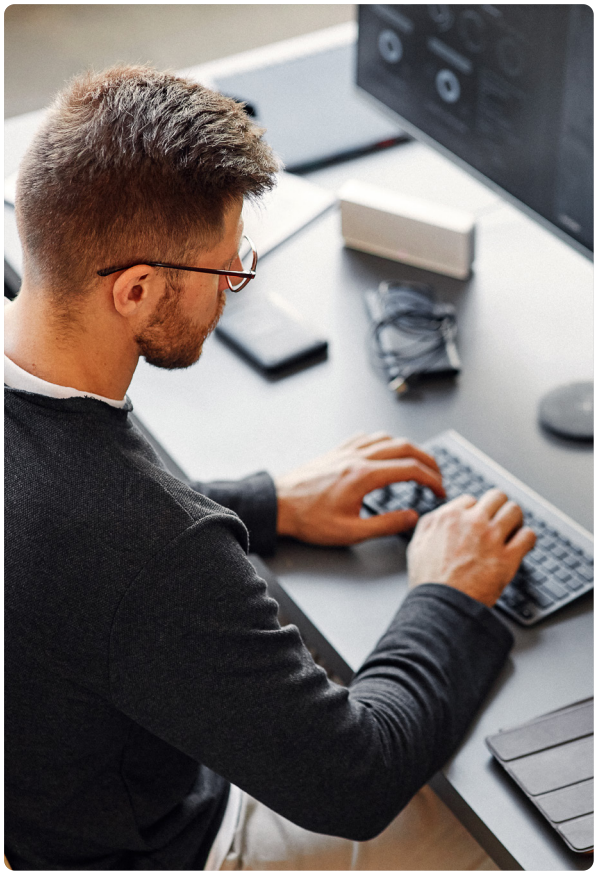

Organizations typically progress through four organizational designs:

- **Centralized AI organization**
  A single central team owns AI delivery and use-case prioritization. Business units act as domain experts and consumers of AI solutions. This works in early stages when AI expertise is scarce and governance needs to be tight, but it concentrates authority where the people doing the work are not.

- **Supported AI organization**
  A central AI hub develops use cases; business units own lifecycle management and value realization. Accountability begins to distribute while expertise remains centered.

- **Federated AI organization**
  Analytics and AI teams sit within business domains and manage use cases end-to-end. A lean central function provides governance, standards, and talent development. Ownership is real, but the AI work and the operational work are still organizationally separate.

- **Product-led AI organization**
  Cross-functional product teams own AI products end-to-end across their lifecycle. A central hub provides platforms, standards, and governance. The people who work alongside the agents also shape them. New roles emerge specifically for judgment, trust, and supervisory oversight of agents.

The key design choice is how responsibility for value creation, delivery, governance are distributed across the enterprise for effective judgement, trust & supervisory of AI

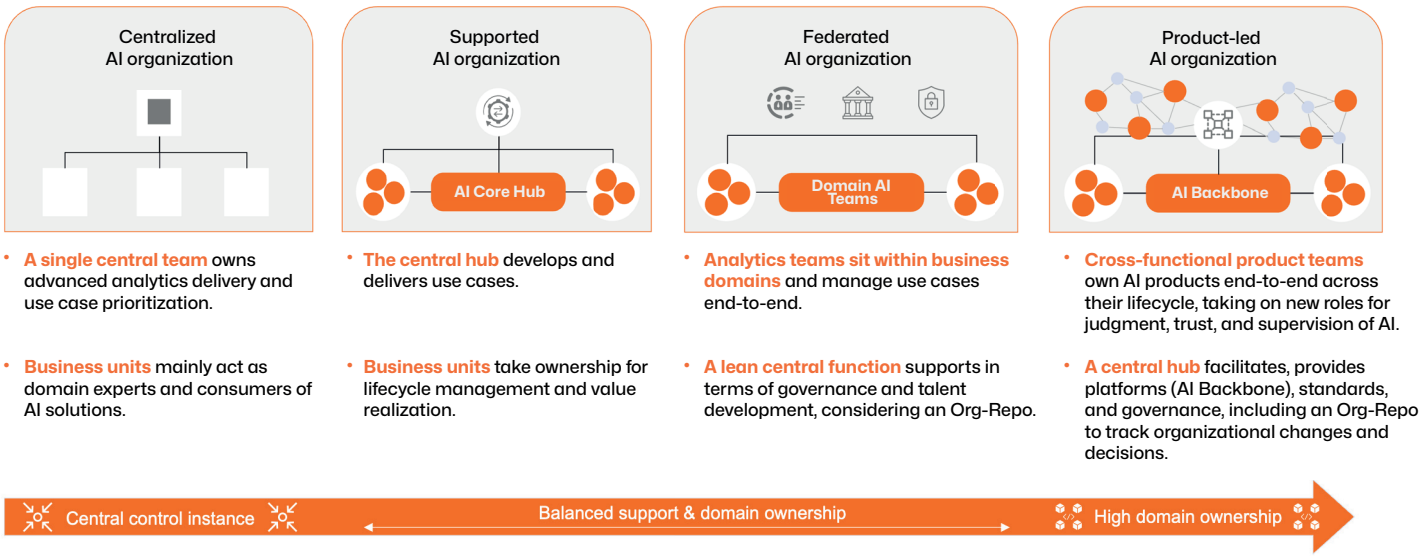

Figure 3: The four archetypes plotted on a continuum from central control to high domain ownership

For most organizations on the path to becoming fully AI-empowered, the **product-led model is the most compelling end state.** It is also the only archetype that is structurally consistent with the bottom-up trust thesis this paper has set out. Co-design and active participation are not occasional interventions in a product-led organization; they are how the work happens by default. Humans and AI agents collaborate within defined product teams, each with clear ownership, shared goals, and built-in guardrails.

This is the structure best suited to the Scale phase, where agent-native processes and headcount-independent growth become achievable – and where the methodology described in the previous section becomes self-reinforcing rather than program-dependent.

The shift to a product-led model does not happen overnight. Organizations typically start with more centralized approaches and evolve as AI capabilities, governance structures, and confidence in working alongside agents all develop together. The AI Way of Working is what makes that evolution possible: it is the bridge from the centralized state most organizations start in to the product-led state where AI value compounds without the central function having to push it.

Making it work: Principles for the journey

Organizations that successfully translate executive ambition into frontline adoption, and progress through the transformation arc, operate by a small number of principles.

→ **Start with the confidence gap, not the use case**
Identify the function where the gap between executive confidence and frontline confidence in AI is widest. That is where the next pilot should run, not where the technical case is easiest. The pilot's purpose is to close the gap visibly, not just to prove the technology.

→ **Active participation beats passive training**
Frontline workers in organizations relying on occasional training programs alone report sharply lower confidence in AI than those in organizations where AI is integrated into everyday work, where role-specific examples are available, and where time has been allocated for hands-on experimentation. Formal training as a standalone intervention is one of the weakest interventions we have measured. Build the program around participation and applied use, not classroom hours.

→ **Match oversight to risk, deliberately**
Full agent autonomy is not the default for every task. Different activities require different authority levels: closely supervised in some cases, distantly monitored in others. The choice should be explicit and reviewed regularly, not inherited from an off-the-shelf governance template.

→ **Treat agents as team members, not tools**
Organizations that reach Reinvent and Scale do not refer to agents as „the AI." They refer to them by name, with defined roles and consistent behaviors. This linguistic shift sounds trivial; it is structural. It signals, and shapes, how authority is shared.

→ **Build use cases and underlying capability in parallel**
Organizations that wait for „the data to be ready" before introducing AI never start. Organizations that ignore data and architecture for too long stall in pilot purgatory. The right pattern is concurrent investment, with each use case strengthening the technical and organizational layers it depends on.

→ **Treat the operating-model redesign as the prize**
Productivity gains from individual AI tools are real but limited. The compounding return, and the competitive moat, comes from redesigning the operating model so humans and agents share authority by default. That is what the journey is for.

Where to start: the Eraneos AI Maturity Diagnostic

The right starting point for any organization depends on where it sits today across the eight building blocks, what its industry constraints are, and how the gap between executive ambition and frontline confidence presents in its specific context.

The Eraneos AI Maturity Diagnostic is the entry point we use with new clients. It produces three things in a four-week engagement:

1. **A position map.** Where the organization sits today across the eight building blocks, benchmarked against peers in its industry and against the Eraneos People & AI Study cohort.

2. **A trust-divide cut.** The size and shape of the gap between executive and frontline confidence in the organization, and where it is widest.

3. **A first-90-days plan.** A focused set of moves that close the most economically important confidence gap and deliver a first measurable Augment-phase outcome.

The diagnostic does not commit the client to a long program. It produces a defensible view of the starting point and the most valuable next move. Most clients then choose a focused 90-day pilot in the function where the diagnostic identifies the highest combined value-and-trust opportunity.

Request your AI Maturity Diagnostic
Contact one of the authors below. We will respond within 48 hours with a scoping conversation tailored to your industry.

Get in touch

At Eraneos, we understand that AI transformation is not just about technology. It is about rethinking how organizations create value and how trust is built between people and the systems they work alongside. Our approach is grounded in four commitments:

Leadership alignment
We work with CEOs and executive teams to connect AI ambition to measurable business outcomes, turning strategy into an executable program.

Bottom-up trust building
We treat the distance between executive ambition and frontline adoption as the central design problem of AI transformation. Our People & AI Study informs every engagement, and our enablement work is designed to close that distance measurably.

AI-native process design
We redesign processes for the agentic era, not by layering AI onto legacy workflows, but by building new workflows designed from the ground up for human-agent collaboration, with the agentic-era risks engineered in from day one.

Readiness for scale
We help organizations build the data foundations, architecture, governance, and trust infrastructure that allow AI to move from pilot to production to enterprise-wide impact.

As a European challenger, we bring a people-first and responsible AI perspective. That means not only accelerating performance today, but building the capabilities, culture, and leadership needed to sustain it. Because real transformation does not happen when AI is introduced. It happens when organizations change how they work.

Contact

Gijs Oostendorp - Associate Partner, AI Transformation
gijs.oostendorp@eraneos.com

Lisa Simon – Senior Director AI First, Strategy
lisa.simon@eraneos.com

Matthias Reck - Principal, Strategy
matthias.reck@eraneos.com

Dr. Olaf Radant - Principal, Organizational Excellence & Transformation
olaf.radant@eraneos.com