HIGHBERG



A positioning paper by Highberg on the organizational prerequisites for successful AI adoption at scale

Designing for AI

The Building Blocks of a Future-Proof AI Operating Model

Introduction: The Organizational Problem Behind the AI Promise

Organizations are investing billions in artificial intelligence today. Budgets are being allocated, talented people recruited, pilots launched. And yet: the majority of AI initiatives never reach production. Models are built but not used. Experiments generate insights but not scale. Boardroom enthusiasm collides with operational friction.

The vast majority of AI investments fail to deliver measurable business impact. Not because of poor technology, but because of poor organizational design.

AI rarely fails in the technology. It fails in governance, ownership, strategic focus, and the ability to move pilots into production. But there is a second, even more fundamental cause that is almost always underestimated: people do not change their behavior and leaders have not been equipped to facilitate that change. Decision-makers do not trust models. Employees continue working the old way. Leaders still rely on intuition rather than data. And even when people want to change, the organization's structures, job descriptions, performance metrics, incentive systems, continue rewarding the old way of working. This is not solely a technical problem. It is an organizational and behavioral problem.

Why This Matters Now

AI is rapidly shifting from "experimental technology" to a core component of business operations. This has three immediate consequences. AI solutions are becoming part of primary processes, requiring uptime, controls, ownership and audits. Risks are becoming more visible: bias, explainability, security, and privacy are no longer hypothetical concerns but concrete boardroom issues. And value is only realized through adoption and scale, which requires organizational design and behavioral change, not yet another pilot.

Bottom line: those who want AI to deliver returns must not only build or deploy AI systems but build the organization and the people capable of sustaining AI.

What Most Organizations Get Wrong

After guiding dozens of organizations, we consistently see the same patterns. Four mistakes recur time and again:

1. They start from technology, not from value. The question is "what can we build?" rather than "where is the value?" The result is a portfolio full of technically interesting experiments with limited business impact and an ROI that nobody can justify.

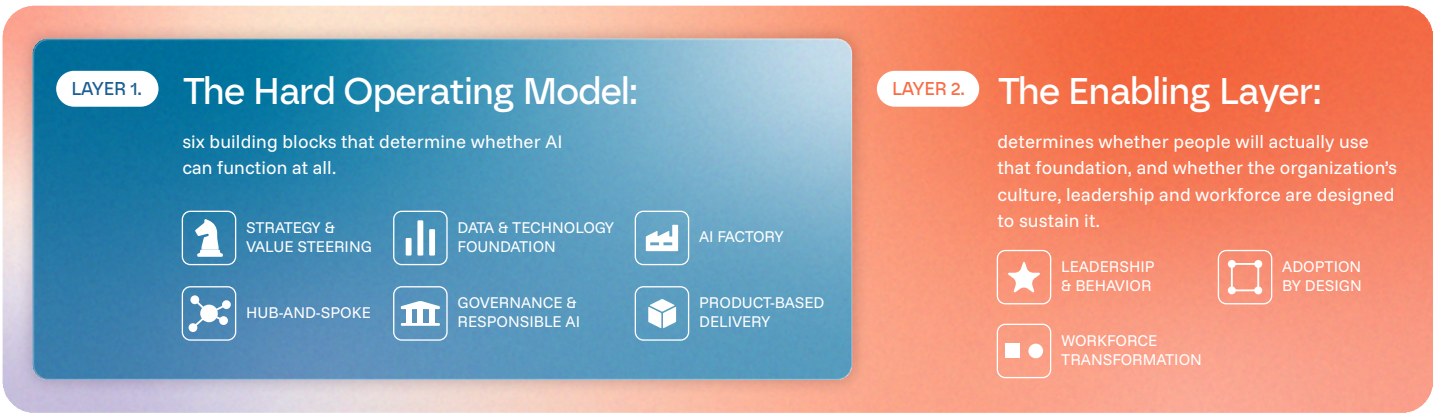
2. They treat AI as a project, not a product. There is a delivery, a go-live, and then it is finished. But AI systems degrade the moment nobody is responsible anymore. Models deteriorate, users disengage, and the investment quietly evaporates.

3. They treat data as an IT problem, not a business asset. Data ownership sits with the department that happens to manage the systems, not with the people who understand the context. The result: AI teams spending 80% of their time on data preparation rather than value creation.

4. They forget that technology does not change people automatically. Governance and structure are necessary but not sufficient. The organizations that truly achieve scale invest just as hard in leadership, behavior, adoption and workforce transformation as they do in models and platforms.

The Highberg AI Operating Model

In this paper we introduce the Highberg AI Operating Model. It consists of two interconnected layers, each answering a different but equally essential question:



Layer 1.
The Hard Operating Model: six building blocks that determine whether AI can work

- Strategy and Value Steering: directing AI toward value, not toward possibilities
- Data & Technology Foundation: the foundation that supports everything
- AI Factory: the engine for repeatable, scalable AI development
- Hub-and-Spoke: central governance combined with decentralized speed
- Governance & Responsible AI: decision structures that enable as a competitive advantage
- Product-based Delivery: AI as a lasting value creator

Layer 2.
The Enabling Layer: behavior, leadership, adoption and workforce transformation determine whether AI will actually work

The first layer gives the organization the structure to build and manage AI. The second layer determines whether people actually use that structure, and whether the organization is built to sustain it. Both are indispensable. Organizations that invest only in the first layer build a perfect machine that nobody uses, operated by a workforce that has not been designed for the new reality.

The Highberg AI Operating Model Maturity Framework

Where does your organization stand, and what is the urgency to move forward?

Level 1. Experiment-driven
Pilots emerge bottom-up, driven by individual enthusiasm. No shared standards, no structural ownership, no repeatable path from experiment to production. Leaders talk about AI but do not steer on it.

Risk: money burns without demonstrable impact. Without intentional direction, AI remains a hobby rather than a capability.
The move to Level 2 requires one thing above all: organizational intent. Someone must decide that AI deserves a structural approach, even if that approach is still local.

Level 2. Department driven
The intent is there, but AI is picked up department by department. Local successes emerge, but without shared standards or central governance. Multiple islands develop in parallel, each with their own tools and logic.

Risk: fragmentation that becomes increasingly expensive to undo. Technical debt accumulates, duplication multiplies, and AI initiatives end up unable to talk to each other.
The move to Level 3 is the most politically demanding transition. It requires departments to give up autonomy in favor of organizational coherence, and leadership at CEO or COO level to make that call.

Level 3. Structured
AI is approached as an organizational capability. Central governance, shared standards and a hub-and-spoke structure are in place. Models reach production in a repeatable, standardized way. But the connection to strategic value is still developing, adoption is managed as a project, and leadership has not yet fully modeled the behavior it expects from others.
Risk: a structured but underperforming organization. The machinery is built, but the returns fall short of what is possible.

The move to Level 4 is less about new structures and more about deepening what is already there: steering deliberately on value, embedding adoption as a design principle, and making leadership behavior the decisive variable.

Level 4. Scaled
AI is fully integrated into how the organization creates value and makes decisions. Product-based delivery is the norm, value steering drives the portfolio, and data is treated as a strategic asset. Critically, the Enabling Layer is active: leaders decide based on data, behavioral change is designed not hoped for, and AI capability is built structurally across the organization.

Advantage: structural competitive advantage. The question is no longer whether AI works, it is how fast you can move to the next frontier.

Most organizations we work with sit between Level 2 and Level 3. The leap to Level 3 is technical, organizational and human.

Building Block 1. Strategy and Value Steering

AI only delivers value when you explicitly choose where to invest and where not to. Without steering, it remains expensive experimentation.

What Is It?

Strategy and Value Steering is the building block that determines what the organization does with AI, and what it does not. It is about translating the AI strategy into concrete value pools and naming explicitly which type of value each initiative pursues. For instance cost reduction, revenue growth or market penetration, customer experience, or risk reduction: where does the real business value lie, which use cases are prioritized, how are investments tracked, and when is an initiative stopped or scaled?

Most organizations have an AI strategy on paper. What they lack is a systematic mechanism to translate that strategy into daily decisions about priorities, budgets and capacity.

Why It Matters

AI investments only deliver value when organizations explicitly choose where to focus. Without clear steering, a fragmented landscape of initiatives emerges, lacking coherence, priority or measurable impact.

For executives, this means that AI without clear value steering leads to high investments with unclear ROI. The absence of prioritization results in fragmentation, delays and loss of focus. And there is insufficient grip on where AI actually contributes to strategic objectives.

This speaks directly to the core questions in every boardroom: where does AI create visible value, where should we stop, and where should we accelerate? Value Steering transforms AI from a collection of experiments into a managed investment portfolio, where choices are explicit and value is continuously monitored and adjusted.

Why Does This Fail in Practice?

The most common cause of disappointing AI results is not poor execution, it is the wrong choices or no choices at all. Often compounded by the wrong expectations and promises of easy cost savings or a fixed percentage of headcount reduction that the technology was never going to deliver on that timeline. Organizations launch use cases based on technical feasibility or internal enthusiasm rather than strategic value potential. The result: a portfolio full of interesting experiments with limited business impact.

Typical signs

- → Use case selection based on available data rather than business value
- → No mechanism to stop underperforming initiatives
- → Benefit tracking is absent, successes are claimed but not measured

What Works: Design Principles

An effective Value Steering mechanism has three elements. First, a structured demand funnel: a clear process for submitting, evaluating and prioritizing AI use cases based on strategic fit, value potential and feasibility. Second, a value case per initiative: no model

goes into production without a measurable hypothesis about what it will deliver. Third, a portfolio governance rhythm: periodic reviews of the AI portfolio to determine which initiatives should scale, which should be stopped, and which deserve more investment.

The Biggest Challenge: From Ambition to Choice

The most fundamental challenge in Value Steering is the willingness to say no. Every department has a compelling AI use case. The political pressure to serve everyone is significant. But organizations that dabble in everything achieve scale in nothing. Leaders who break through this pattern treat their AI portfolio as an investment portfolio: with deliberate allocation, regular reassessment and the discipline to disinvest when results fail to materialize.

△ If you get this wrong

- Your AI portfolio grows in breadth but never in depth: dozens of initiatives, none at scale
- Budgets are extended year after year without demonstrable business impact
- AI loses credibility in the boardroom before it ever demonstrates its potential

→ First steps

- Define 3 to 5 value pools: where do money, risk or customer impact reside?
- Establish a portfolio board with a regular stop-scale rhythm, for example monthly
- Make benefits tracking mandatory for every AI initiative

🗣 Key roles

Chief AI Officer / AI Strategy Lead. Translates business strategy into an AI agenda with concrete priorities and measurable ambitions.

Value Manager / Business Case Owner. Monitors value hypotheses per initiative and ensures benefit realization is actually tracked and reported.

Portfolio Governance Board. The decision-making body that periodically reviews the AI portfolio and adjusts based on value and strategic fit.

Building Block 2. Data & Technology Foundation

Without reliable data & technology and clear ownership of it, every AI investment will underperform.

What Is It?

Data is not a subcomponent of AI, it is the foundation on which everything rests. Yet most organizations treat data as a technical byproduct: something the IT department handles when needed. This is a fundamental misconception that leads to structural failure.

A strong Data Foundation encompasses three things: the technical availability and quality of data, organizational ownership of data per domain, and treating data as a strategic asset with its own lifecycle management, quality standards and governance.

Why It Matters

Without reliable and accessible data, AI remains a cost center rather than a value accelerator. Investments in models deliver higher returns when the underlying data is accurate, available and owned by the business.

For executives, this means data is not an IT issue but a business asset with direct impact on risk, compliance and performance. A lack of data ownership translates directly into delays, higher costs and unreliable outcomes. Organizations that manage data well achieve scale faster and realize higher ROI on AI investments.

In short: without a data foundation, every euro invested in AI is suboptimally spent.

Why Does This Fail in Practice?

In virtually every AI initiative that stalls, data plays a central role. Not because data is absent, organizations typically sit in a sea of data. But because that data is unavailable, unreliable or misunderstood at the moment it is needed. Data ownership is the most structural problem: in most organizations, nobody is specifically responsible, which means AI teams spend 70 to 80% of their time on data preparation rather than value creation.

Typical signs

- → Data ownership sits with IT rather than with the business that understands the context
- → No data quality standards, every team works with its own definitions
- → Data treated as a project asset rather than an organizational asset

What Works: Design Principles

Data product thinking is the guiding principle. Data is not treated as a raw material used once, but as a product: documented, equipped with quality standards, actively maintained and made available to multiple users. Ownership lies with the business, not with IT. IT facilitates and supports, but the business domains are accountable for the quality and accessibility of their data.

The Biggest Challenge: Embedding Data Ownership in the Business

The technical challenges around data are solvable. The organizational challenges are too, but they are more persistent. Business managers experience data ownership as an additional burden on top of their primary responsibilities, until an AI model demonstrably gives poor advice due to poor data. Organizations that handle this well link data ownership to business results: data quality becomes part of the performance measurement of the domain.

⚠ If you get this wrong

- AI teams lose the majority of their capacity to cleaning data rather than creating value
- Models are built on unreliable foundations, producing incorrect outcomes and reputational damage
- Data incidents and compliance failures are not a matter of if, but of when

→ First steps

- Appoint data owners for the top 3 critical domains
- Start with 10 core datasets: establish definitions and quality metrics
- Implement a minimal metadata and lineage approach for auditability

🗣 Key roles

Chief Data Officer / Data Governance Lead. Ultimately responsible for the data strategy and data governance framework of the organization.
Data Domain Owner (per business domain). The business owner accountable for data quality and accessibility within their domain. This is an explicit accountability, not an informal role.

Data Engineer. Builds the pipelines and infrastructure that make data available, repeatable and reliable for AI applications.

Data Steward. Operationally responsible for data quality within a domain: signals deviations, monitors definitions and ensures alignment with central standards.

Building Block 3. AI Factory

Without industrialization, AI remains stuck in pilots. Scalable value requires a repeatable delivery model.

What Is It?

The AI Factory is the standardized, repeatable way in which an organization develops, tests and brings AI solutions into production. Most AI investments do not fail, they stall. Not in the technology, but in the absence of an operational delivery engine that makes the transition from experiment to production structural and repeatable.

The AI Factory does not think in isolated projects but in a production process: fixed steps, reusable components, predictable timelines. From data extraction and model development to deployment and monitoring, supported by shared platforms and tooling.

Why It Matters

Most AI investments stall between experiment and production. This is not a technical problem but an operational one: a scalable way of building and managing AI solutions is absent.

For executives, this means that without industrialization, AI remains stuck in pilots with a low ROI on the total investment. Time to market stays unpredictable and slow. And the absence of standards leads to fragmentation and rising costs. An AI Factory transforms AI from experiment into a predictable value engine, which is essential for scaling AI and realizing competitive advantage.

*An Explicit Note on Generative AI

Generative AI is changing the rules of the AI Factory in three ways. The speed of use case development increases dramatically, what previously took months can now be realized in weeks. At the same time, LLM-based applications introduce new risks: hallucinations, prompt injection and unintended data leakage via model outputs. And the type of use cases is shifting: from predictive models toward knowledge work, decision support and automated agent workflows. A modern AI Factory is equipped for both worlds, classical ML and generative AI, with the corresponding standards, evaluation frameworks and governance gates for each type of application.

Why Does This Fail in Practice?

Every team or employee builds in its own way, with its own tools, its own definitions and its own assumptions. Deployment is an afterthought: the organization invests heavily in model development but hardly at all in the infrastructure needed to get models reliably into production. The result is a paradox: a great deal of AI activity with a poor ROI.

Typical signs

- → Too much focus on model development, too little on deployment and monitoring
- → Tooling without standardization, every team chooses its own stack
- → Data as a structural bottleneck that is never fundamentally addressed

What Works: Design Principles

The AI Factory operates as a production line: every step is defined, every output is standardized and reusable components are shared. What good looks like: 80% of new AI use cases follow the same standardized pipeline, and the average time from pilot to production is less than four weeks. That is ambitious, but for organizations at Level 4 it is the reality.

The Biggest Challenge: Industrialization as an Organizational Culture

The deepest challenge of the AI Factory is cultural. Data scientists want to build and experiment. Standardization feels like a constraint. The AI Factory requires one fundamental shift in mindset: from "we build the best model" to "we deliver working AI into production, repeatedly and predictably." Leaders who handle this well treat deployment as a first-class concern, not a detail to be addressed later.

△ If you get this wrong

- 80% of your AI investments remain stuck in pilots: visible to the organization but without impact
- Costs rise without economies of scale: every use case is reinvented from scratch
- AI loses credibility as something that "always promises but never delivers"

→ First steps

- Select 1 to 2 reference use cases and industrialize them end to end
- Define minimum standards: logging, monitoring and model cards
- Position MLOps as a must-have, not a nice-to-have

🗣 Key roles

ML Engineer / MLOps Engineer. The link between experiment and production. Automates deployment, monitors model performance and ensures models are maintainable. Without this role, the best models stall in the pilot phase.

Data Scientist. Builds and trains models, selects algorithms and validates whether a solution actually answers the business question.
AI / LLM Specialist. With generative AI, the focus shifts from "building models" to "deploying models intelligently." Prompt engineering, RAG architectures and designing AI agents require specific expertise.

Platform Engineer. Manages the shared AI infrastructure: compute, storage, experiment tracking and model registry. Prevents teams from building their own infrastructure independently.

Building Block 4. Hub-and-Spoke

The balance between central governance and decentralized execution determines whether AI accelerates or stalls.

What Is It?

The hub-and-spoke structure resolves a classic organizational dilemma: centralize AI and you lose speed and connection to the business. Decentralize AI and you lose consistency, efficiency and control. The organizations that handle this best do not choose one or the other, they design the boundary deliberately.

The hub is a central AI capability, a team or center of excellence, responsible for standards, tooling, platforms and knowledge sharing. The spokes are the business domains that, supported by that central infrastructure, apply AI with their own priorities and pace.

Why It Matters

AI touches virtually every part of the organization. Without a clear definition of central and decentralized responsibilities, the result is either sluggishness or chaos.

For executives, this means that a poorly designed structure leads to a loss of speed or a loss of control. Ambiguity about mandate and funding creates internal friction and political disputes. And AI becomes trapped in silos, meaning the benefits of scale are never realized. The right balance between central and decentral determines whether AI becomes an organization-wide capability or remains a fragmented initiative.

Why Does This Fail in Practice?

Without explicit rules, two destructive patterns emerge. In the first pattern, the hub becomes an IT bottleneck: business teams go their own way, resulting in fragmentation and inconsistent tooling. In the second pattern, everything remains centralized, but AI loses its relevance because it is too far removed from the business reality.

Typical signs

- Unclear division of responsibilities: who is authorized to decide what?
- A central hub that thinks from technology rather than from business value
- Spokes that embrace AI without respecting the central standards

What Works: Design Principles

The key is an explicit governance document, not as a bureaucratic instrument, but as organizational design: who is accountable for what, which decisions are central and which are decentralized, and how are conflicts resolved? As a domain becomes more capable, more ownership shifts to the spoke. What remains central are the rules of engagement, not the execution.

The Biggest Challenge: Politics and Rules of Engagement

Hub-and-spoke is the most politically sensitive design choice in an AI operating model. Both sides of the frustration are understandable and both are resolvable, but only if the rules of engagement have been made explicit and are anchored at the level of the CEO or COO, not just the CISO or CDO.

⚠ If you get this wrong

- Business teams build their own AI islands: duplicate work, inconsistent outcomes and unmanageable risks
- The central hub becomes a brake rather than an accelerator, and loses legitimacy
- The economies of scale from shared platforms and knowledge sharing are never realized

→ First steps

- Document RACI and funding arrangements: one page, not a lengthy report
- Start with 2 spokes: high potential and genuine willingness
- Make the hub output tangible: templates, patterns and platform services

🗣 Key roles

Head of AI / Chief AI Officer. Directs the central AI ambition, maintains the coherence of the model and acts as the strategic counterpart for the C-suite.

AI Translator / Analytics Translator. Connects business questions to technical reality. Translates what a director wants to know into what a data scientist can build, and vice versa. Without this role, business and technology consistently talk past each other.

Business Product Owner (in the spoke). The local driver of AI value. Responsible for the prioritization and adoption of AI solutions within their domain. This is a business role, not an IT role.

Platform Owner (in the hub). Manages the shared tooling and infrastructure. Prevents technical fragmentation and maintains the consistency of the AI foundation.

Building Block 5. Governance & Responsible AI

AI does not require more control, but clearer decision structures: who decides, based on what values, with what accountability. Governance is the mechanism, responsible outcomes are the goal.

What Is It?

Governance is the least glamorous part of an AI operating model, and therefore the most neglected. Until something goes wrong. A model that disadvantages certain groups in credit decisions. A generative AI tool that leaks confidential information. An automated decision system that cannot explain why it made a particular decision. The reputational damage and legal consequences are substantial, and they are preventable.

But effective governance is not primarily about preventing harm: it is about designing responsibility in from the start. Organizations that integrate privacy considerations, security principles and AI risk assessments into their development process from the beginning do not just avoid incidents. They build AI that is more robust, more trustworthy and easier to defend: to regulators, to users and to the board. Compliance becomes a quality standard rather than a checkpoint. This is what responsible AI by design means in practice: not reactive oversight, but a discipline embedded into every stage of the AI lifecycle, from use case selection through to production.

Why It Matters

AI introduces new risks that are directly visible at board level: bias, compliance, reputation, security and explainability. Incidents have not only operational but also strategic and reputational impact.

For executives, this means well-designed governance structures, those that embed the right values, accountability mechanisms and oversight principles from the outset, are the operational precondition for responsible AI. Governance by itself does not create trust; responsible decision-making applied through governance does. Poor governance can lead to direct legal and reputational damage. But overly heavy governance constrains innovation and time to market. The challenge is not more control, but smarter, risk-based control that enables speed.

Why Does This Fail in Practice?

The most common governance mistake is treating governance as an approval process: routing every AI initiative through a series of committees until the momentum is gone. Teams learn to circumvent this. The alternative, no governance, leads to applications going live without adequate checks.

Typical signs

- Compliance, risk and ethical considerations are treated as final approval gates and act as blockers rather than being embedded throughout design, development and deployment.
- No risk-based differentiation: every model subject to the same heavy procedure
- Compliance and AI teams operating as adversaries rather than partners

What Works: Risk Based Governance Structure and Responsible By Design

A risk-based framework is the guiding principle. An internal dashboard has a different risk profile than an automated decision system for healthcare eligibility. Governance effort must be proportional to impact and risk. What good looks like: every AI system has a risk classification, a documentation set and a clear owner, before go-live, not after. Adopt the principle of responsible by design as a way of working.

The Biggest Challenge: Getting Compliance and the Business to join forces.

The deeper governance challenge is cultural: it requires compliance, legal and business teams to see themselves not as adversaries but as joint designers of systems that are both fast and responsible. This requires shared vocabulary, mutual understanding and leadership that takes both dimensions seriously.

△ If you get this wrong

- AI incidents, bias, leakage, incorrect decisions, become public and cause lasting damage to trust
- Legal exposure under the EU AI Act increases as AI becomes more deeply embedded in primary processes
- Governance becomes reactive rather than preventive: costly, slow and damaging to the AI program

→ First steps

- Conduct a combined privacy, security and AI risk assessment for your highest-priority use case, before build begins, not after
- Define risk classifications (three is sufficient to start) with proportionate controls for each tier
- Embed responsible AI principles into your use case intake process, so governance shapes decisions from day one rather than reviewing them after the fact

🗣 Key roles

AI Governance Lead / AI Risk Officer. Designs and monitors the governance framework. Translates risk knowledge into practical frameworks that teams can apply without becoming paralyzed.

Legal and Compliance / Privacy Officer. Indispensable now that the EU AI Act and GDPR together introduce obligations that are fundamentally different from what organizations are accustomed to. Early involvement prevents costly revisions.

Model Validator / Independent Reviewer. The independent check on the quality, fairness and reliability of models. A model deemed good internally can produce undesired outcomes in production due to bias, drift or unforeseen edge cases.

Security Specialist. AI works with sensitive data and introduces new attack vectors, particularly with generative AI. Security is not a layer that is added afterwards.

Building Block 6. Product-based Delivery

AI is not a project but a product. Without structural ownership, value disappears after delivery.

What Is It?

The sixth building block is about ownership and continuity. Many organizations treat AI as a project: there is a delivery, a go-live, and then it is done. This approach is fatal, because AI systems degrade the moment nobody is responsible anymore.

An AI model that performs well today will perform worse six months from now. Customer behavior changes. Market conditions shift. The data on which the model was trained represents a reality that no longer exists. Model drift is unavoidable. The question is not whether a model will degrade, but when and whether the organization detects it in time.

Why It Matters

AI only delivers value when it is sustainably used and maintained. Without structural ownership and adoption, value disappears shortly after delivery.

For executives, this means a project-based approach leads to rapid value erosion through model drift and low adoption. Insufficient ownership means nobody is responsible for impact. And investments remain one-off while value requires continuous management. AI is not a project, but an ongoing product that demands structural attention and funding.

Why Does This Fail in Practice?

Budget cycles are structured around projects. Capacity planning is structured around projects. Responsibilities are structured around projects. Once an AI initiative exits the build phase, the urgency disappears, and with it the budget and the talent. The model gradually deteriorates. Users notice, lose confidence and stop using it.

Typical signs

- No clear owner after go-live
- Monitoring as an afterthought: model drift is detected too late
- Adoption as a final item: technical success, but nobody uses it

What Works: Design Principles

AI products are structurally incorporated into portfolio management: with their own budget, their own roadmap and an owner who remains accountable for impact after go-live. What good looks like: every AI system in production has a named Product Owner, a monitoring dashboard, a defined retraining protocol and an annual value evaluation. Adoption is a design principle from the outset, embedded in how the product is built, not in how it is communicated afterwards.

The Biggest Challenge: Ownership as Organizational Design

The deepest challenge in product-based delivery is the willingness to fund AI permanently. The CFO who treats AI budgets as one-off project investments structurally undermines the value of the technology. The shift from project thinking to product thinking is not a detail in the budget, it is a fundamental choice about how the organization defines value. And it requires leaders who are willing to invest political capital to make it happen.

⚠ If you get this wrong

- Models degrade unnoticed: the organization relies on outcomes that have long since stopped being accurate
- Adoption drops to zero after go-live: months of investment evaporate within quarters
- AI becomes internally associated with "interesting experiment, but it never really works", and loses structural support

→ First steps

- Appoint one AI Product Owner for a high-impact use case
- Set up monitoring and adoption metrics from day one
- Incorporate AI products into the regular portfolio and annual plan

🗣 Key roles

AI Product Owner. The central owner of the AI solution across its full lifecycle. Prioritizes the backlog, steers on impact and is the face of the solution toward the business.

Cross-functional AI team. AI requires collaboration that no single discipline can provide alone. A sustainable team combines business, data and engineering, and works continuously on improvement.

Monitoring / Performance Analyst. Monitors model performance in production and signals early when a model deviates. Initiates the process of adjustment or retraining before users notice.

Change and Adoption Lead. Embeds adoption as a design principle from day one. Guides the behavioral change needed for AI to truly land in the organization, and forms the bridge to the second layer of the model.

The Enabling Layer: Behavior, Leadership, Adoption and Workforce Transformation

The first layer determines whether AI can work. The second layer determines whether AI will actually work.

In virtually every AI transformation, the biggest bottleneck is not technology or governance, it is behavior. Decision-makers do not trust models. Employees do not change their working methods. Leaders talk about AI but do not steer on it. The result: AI exists as a system but not as a way of working. It operates as a parallel universe alongside the organization rather than as an integral part of it.

> AI only changes things when people start working differently, not when models get better.

This is the most underestimated truth in AI transformation. And it is the reason why organizations with perfect architectures and solid governance still fail: the enabling layer is missing.

Why It Matters

The six building blocks of the Hard Operating Model create the structure and infrastructure for AI. But structure and infrastructure do not change behavior by itself. Value only emerges at the moment when organizational designs are updated and people, from the boardroom to the shop floor, adjust their decisions, working processes and habits based on what AI offers them.

For executives, this means three things. First: low adoption is direct value destruction. A model that is not used delivers no ROI by definition, regardless of its technical quality. Second: behavior follows leadership. If leaders continue overriding decisions based on intuition, they implicitly give the organization permission to do the same. Third: adoption is not a communications challenge but a design challenge. It must be embedded in processes, incentives and decision-making structures, and ultimately in how the workforce itself is organized and rewarded.

Why Does This Fail in Practice?

Adoption is systematically underestimated and addressed too late. It is treated as the final step in a project, when it should have been the foundation.

- Employees do not trust AI outcomes or accept them blindly and are never guided in building that trust
- Incentives remain unchanged: there is no reason to work differently
- Leaders do not set the right example and thereby undermine the entire program
- Adoption becomes “training” rather than behavioral design: one-off, superficial and quickly forgotten
- Workforce structures, headcount, roles, skills, organizational designs,, are never redesigned to reflect the new reality

What Works: Design Principles

Behavioral change in AI requires a fundamentally different approach from classical change management. Four principles make the difference.

→ **Adoption starts with leadership, not communications.**
Leaders determine through their behavior whether AI works in the organization. If they make decisions based on data and models, the rest will follow. If they do not, AI will die in the organization, regardless of how well the system was built.

→ **Redesign decision-making, not just tooling.**
AI only delivers value when it is embedded in how decisions are made. This means adapting processes, redefining roles & responsibilities, calibrating headcount and deliberately designing human-in-the-loop mechanisms.

→ **Link KPIs to usage.**
Adoption without accountability is voluntary. Organizations that handle this well measure usage, the quality of AI-driven decisions and the impact on business results, and make this visible at management level.

→ **Build capability structurally.**
One-off training does not work. Sustainable adoption requires a structural program of upskilling, AI literacy for management and a deliberate choice between building internally or hiring externally.

The Biggest Challenge: From Convincing to Embedding

The most common mistake in adoption programs is relying on persuasion. People are informed, trained and made enthusiastic, but as soon as the pressure of daily work returns, they fall back into old patterns. Lasting behavioral change requires not only persuasion but embedding: in processes, roles & responsibilities, in systems, in what is measured and rewarded.

> “Adoption does not begin with training. It begins with the right example from management, and it ends only when the new behavior becomes the new norm.”

⚠ If you get this wrong

- The six building blocks deliver a perfect machine that nobody uses
- AI investments fail to deliver returns, not due to technical failure, but due to human failure
- The organization develops AI fatigue: every new initiative meets with skepticism because previous initiatives delivered nothing

→ First steps

- Make leaders explicitly accountable for AI adoption in their domain, not as an HR matter but as a business priority
- Redesign one critical decision-making process with AI as an integral component and adjust FTE, roles, responsibilties and hierarchy to reflect the new situation
- Launch an AI literacy program for management, not technical, but focused on decision-making and trust

🗣 Key roles

Business Leadership (primary). The most crucial and most underestimated role in the enabling layer. Leaders determine through their behavior whether AI works in the organization. This is not a responsibility that can be delegated.

Change and Adoption Lead. Designs and guides the behavioral change. Not as a project manager of a communications campaign, but as an architect of new ways of working and decision-making processes.

AI Translator. Serves alongside the technical bridging function also as a behavioral bridge: helps employees and managers understand what AI can do, what it cannot do, and how to work with it effectively.

HR / Learning and Development. Responsible for the structural development of AI capability in the organization: upskilling, reskilling, AI literacy and a long-term hiring and development strategy.

Workforce Transformation Lead. Designs and drives the structural redesign of roles, functions and reward structures. Ensures that the organizational design reinforces the behaviors the AI program requires, rather than working against them.

Conclusion: Structure, Behavior and Workforce Design Determine the Difference

The six building blocks of the Hard Operating Model give the organization the structure to build, manage and scale AI. The Enabling Layer determines whether that structure actually comes to life, and whether the organization's people are designed, developed and rewarded in a way that makes AI sustainable.

Together they form an integral transformation perspective and that is precisely what AI transformation requires. Not only a technical or governance challenge, but a fundamental redesign of how an organization works, decides, learns and organizes its people.

Organizations that reach Level 4 have not done so through better algorithms or more data scientists. They have done so by fundamentally redesigning their organization and their people: with explicit value steering, embedded data ownership, a standardized delivery engine, clear division of responsibilities, proportionate governance, sustainable product thinking, and an Enabling Layer that combines the right leadership behaviors with the structural workforce changes needed to make them stick.

The question is no longer whether your organization should embrace AI. The question is whether your organization is structured to make AI work. Structurally, at scale, and in a way that earns the trust of customers, employees and regulators.

That begins with the design of your operating model. And it ends only when behavior and organizational structure change together.

Pick one building block and start this quarter.

Wherever you are today, there is a concrete first move. We help you:

- → Identify the one building block where action now creates the most value
- → Define a small, concrete first step you can take in the next few weeks
- → Build a concrete 90-day roadmap with stop, start and scale decisions

Get in touch for a no-obligation conversation

Ralph van den Bosch
Partner
+31 (0)6 535 018 35
ralph.vandenbosch@highberg.com

Gido van Puijenbroek
Partner
+31 (0)6 159 155 90
gido.vanpuijenbroek@highberg.com

Laura Natrop
Partner
+31 (0)6 207 035 92
laura.natrop@highberg.com

This position paper reflects the collective insights and experience of the authors.

AI tools were used as a co-creation aid in drafting and refining the text; all conclusions, judgments, and final content are the authors' own.

Assisted by AI