LastingImpact.AI
your future is in your data

WHITEPAPER

The Hybrid AI Innovation Model

How to balance bottom-up and top-down AI innovation, and make the progress legible to your board.

LastingImpact.AI

May 2026, revised edition

Executive Summary

Most enterprise AI programs underdeliver, and the cause is rarely the technology, but more often a design flaw. Each picks one of two incomplete architectures, top-down redesign by a central AI team or bottom-up experimentation by frontline employees, and each fails in its own predictable way. The cost of staying with a broken architecture compounds: McKinsey's high performers are already three times more likely than their peers to have redesigned workflows around AI.[4] And Article 4 of the EU AI Act already makes AI literacy a legal obligation for staff who work with AI; national enforcement begins on 2 August 2026, which gives boards a fixed date to act against.[15]

Each of the two architectures has real evidence behind it. Bottom-up innovation adopts faster, fits operational reality, and earns stronger ROI on the individual use cases.[1] [3] [5] But without a coordinating architecture, distributed experiments rarely add up to real, AI-first transformation. Top-down redesign, on the other hand, supplies the mandate and architectural coherence that transformation needs.[4] [6] [7] Designed away from the floor, though, it stalls on adoption and feasibility. Even PwC, bottom-up's sharpest critic, concedes that an orchestration layer can "help end-user innovation enhance your top-down strategy."[8] That concession describes the hybrid model in all but name.

This is the paper's argument: each strategy fails alone, and the evidence supporting each is exactly why neither can be dropped. The Hybrid Approach joins the two streams through a shared governance mechanism: a VC-style innovation funnel feeding a hub-and-spoke AI Centre of Excellence.[12] It is a model we have delivered. In one engagement it lifted a service line's productivity by 50%; in another it cut a portfolio's unnecessary field visits by 15–18% while building a capability the business kept (Section 6).

The Core Argument. Bottom-up innovation produces feasible, well-adopted use cases anchored in operational reality.[1] [3] [5] Top-down strategy supplies the mandate, vision, and architectural coherence transformation requires.[4] [6] [7] [8] Neither works without the other. Governed like a VC portfolio and structured as a hub-and-spoke Centre of Excellence,[12] the Hybrid Approach makes the two mutually reinforcing.

1. The Evidence Base

The research behind each approach is substantial. Here are the most-cited 2025–2026 findings, set side by side.

The case for bottom-up.

MIT NANDA's 2025 study found that 95% of enterprise generative-AI pilots never reach production. The organisations that cross the divide empower line managers rather than central labs.[1] Johnson & Johnson's experience shows both halves: 900 use cases sourced from the frontline, but only 16–17% converted to production until decision rights moved to domain-level leaders.[2] IMD finds reskilling pays off only where infrastructure and motivation already align.[3] And Wharton's survey of more than 800 leaders is the counterweight to the failure narrative: 82% use generative AI weekly, and around 75% report positive (self-reported) ROI.[5]

The case for top-down.

McKinsey identifies just 6% of organisations as AI high performers, firms three times more likely than peers to have redesigned workflows around AI,[4] and it puts a number on the human side: five dollars should go to people for every dollar spent on technology.[6] BCG finds 72% of CEOs now act as their company's principal AI decision-maker, with Trailblazers directing 60% of AI budgets to upskilling against 27% for laggards.[7] Deloitte ranks insufficient skills as the single biggest obstacle to AI value, ahead of technology and budget.[9] PwC presses hardest: ungoverned crowdsourcing yields misaligned, duplicative, unscaleable projects.[8] And HBR locates the binding constraint at the "last mile," where capability has to become operational reality.[11]

Both bodies of evidence are real, and both are incomplete alone. The model that honours both is the one integrating them together, described in Section 3 below.

2. Two Innovation Modes: Strengths and Limits

Before combining the two, it helps to see what each does well and where each reliably breaks.

The bottom-up approach. Business experts, trained in AI fundamentals, ideate and prototype with no-code tools. Its strengths come from where the work starts. Use cases stay close to operational reality. The expert who proposes a solution is usually the one who will use it, so adoption is built in. And the organisation gains a cadre of AI-literate employees who can act as informed clients afterwards. It also carries a built-in limit: frontline innovators optimise within the structures they already know rather than redesigning them, which produces the patchwork PwC describes when no governance sits above it.[8]

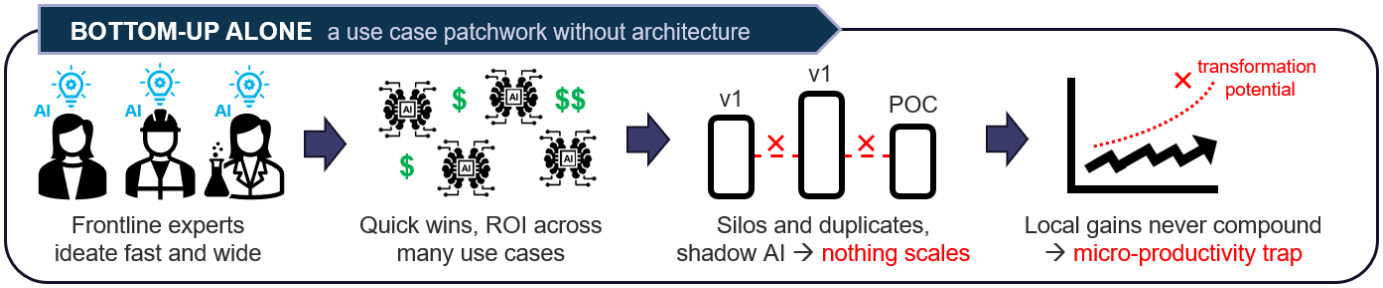

The top-down approach. A central team designs how a domain should look once AI is fully embedded: benchmarks, human-machine task allocation, and the capabilities the design will need. Its strengths are ambition, coherence, and the mandate to make structural decisions. Its risks trace to distance from the floor: use cases designed without operational knowledge, teams too passive and too little AI-literate to challenge them, critical-path dependencies no one has prototyped, and adoption resistance that surfaces too late, much of it concentrated at the mid-management layer.[10] These programs typically take 18–36 months to show results,[11] long enough for sponsorship to erode first.

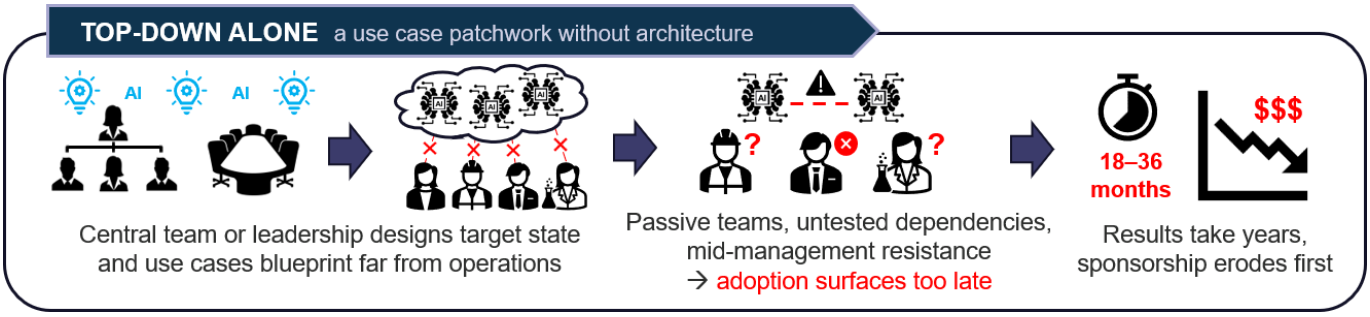

The table below sets the two against each other, and against the hybrid that resolves their opposing weaknesses.

Dimension	Bottom-Up Only	Top-Down Only	Hybrid Approach
Mandate & Vision	Organic; fragmented without coordination	Strong mandate to transform end-to-end	Top-down vision; bottom-up value realisation
Idea Generation	Many ideas from frontline; fast and wide	Small strategy team; ivory-tower risk	Roadmap-guided; populated with validated ideas
AI Literacy	Employees upskilled firsthand	Stays with consultants or CoE; teams passive	Broad bottom-up literacy; shared co-design language
Prototyping & Validation	Rapid on real data; feasibility confirmed early	Designed on paper; feasibility untested until late	Roadmap use cases de-risked through expert validation
Product-Market Fit	High; designed by and for the executors	Risk of mismatch; designed without context	High; bottom-up grounded and top-down reality-checked
Adoption Risk	Low; employees champion their own ideas	High; change imposed top-down	Low; business owners co-created the roadmap
Transformation Depth	Shallow; the micro-productivity trap	Deep, when execution lands	Deep; validated, feasible redesign
Cost & Time to Value	Fast quick wins; high ROI per case	Slow, expensive; long delivery	Quick wins fund momentum; roadmap delivers transformation

Dimension	Bottom-Up Only	Top-Down Only	Hybrid Approach
Fragmentation Risk	High; a patchwork of siloed solutions	Low; single integrated design	Low; top-down architectural coherence

3. The Hybrid Approach: A VC-Inspired Innovation Funnel

The Hybrid Approach does more than run both streams at once. It connects them through a shared governance mechanism whose logic comes from a model already proven at managing high-uncertainty portfolios: early-stage venture capital.

Why the VC parallel holds. Venture capital lives in conditions that mirror the enterprise AI problem: many uncertain ideas, high failure rates, and a need to concentrate resources on the few that will return disproportionately. Good VCs invest a little to test many ideas, select hard at each funding gate, and double down only where traction is proven. Every stage points at the exit. For an AI use case, that exit is enterprise-wide deployment of a validated solution. Lose sight of it, and the funnel becomes a prototype factory that never ships.

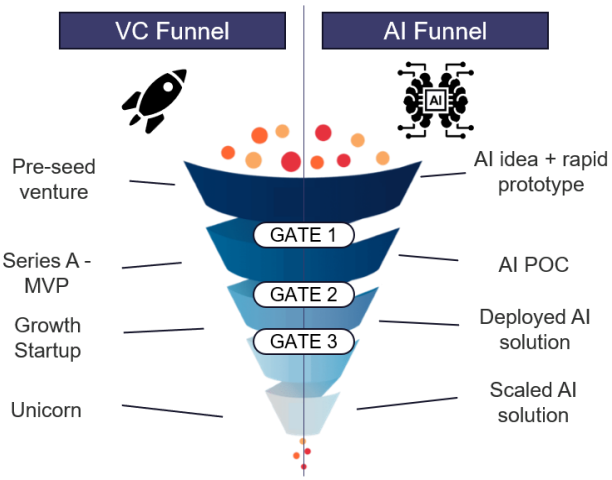

The AI intake: Front line-originated AI use cases ideas and rapid prototypes

AI Governance: Invest a little in many ideas, select hard at each gate, double down on proven traction

The AI 'exit': Company-wide deployment of most successful use cases

For a mid-cap company, governance is what makes the method defensible. It keeps capital from flowing to use cases that never reach production, and it turns validated automation into lasting margin. For the executive accountable for AI, it also produces something a board can read: evidence of progress, a redesigned operating model, and a clear line from spend to outcome. The discipline modernises the company without betting it, and it protects competitiveness and jobs as AI reshapes the work.

There is also a funding mechanism hiding in the cost base. As transformational use cases aggregate, they begin to absorb work previously rented from SaaS vendors, and every subscription retired or reduced converts recurring SaaS licence spend into budget for the next wave of AI use cases. Over time the AI program starts paying for itself out of the software bill it replaces.

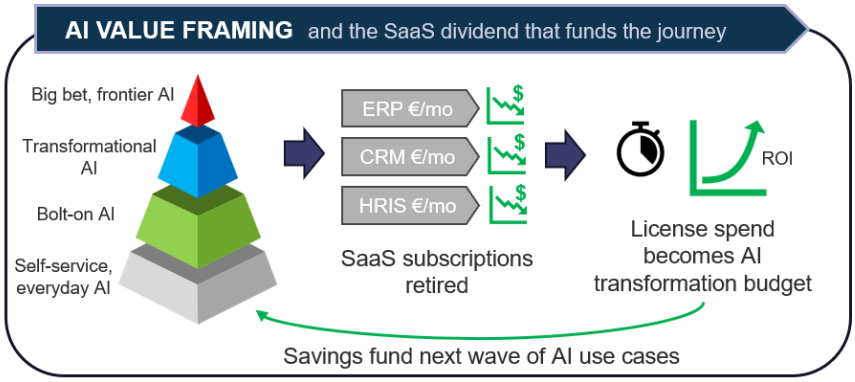

Answering the top-down critique. PwC's charge, that ground-up crowdsourcing produces misaligned, unscaleable projects,[8] is fair against programs with no governance layer. The Hybrid Approach is the correction. A top-down mandate sets where transformation must happen, and a VC-style funnel steers frontline innovation toward those priorities. PwC's own concession names the resolution, and the hub-and-spoke Centre of Excellence is the structure most 2026 practitioner sources adopt to provide it.[12]

The six design principles.

Principle 1: Upskill first, innovate second. Transformation cannot be handed to a technical elite. The first investment is broad AI literacy, usually a two-to-three-day immersive workshop. Article 4 of the EU AI Act already requires it for all staff who work with AI, with national enforcement starting 2 August 2026.[15]

Principle 2: No-code tools to democratise prototyping. Without them, every prototype waits in the AI/IT specialist queue, and the bottleneck returns. A prototype running on real data for two weeks tells you more about feasibility than a design document written over six months.

Principle 3: Generate volume, then filter. The funnel starts wide. Aim for 30–50 use cases per major business domain before consolidating, generated through structured ideation sprints.

Principle 4: Apply VC-style funding gates. Each use case passes defined gates where leadership backs it further, pauses it, or retires it. J&J's pivot to domain-level decision-makers is exactly this move. The output is a de-risked roadmap where investment follows proven potential, not ideas on paper.

Principle 5: Feed validated use cases into the top-down roadmap. Prototyped, real-data-tested, feasibility-scored use cases are mapped against the strategic blueprint in a recurring review. Top-down design meets operational truth, and critical-path use cases are validated before commitment.

Principle 6: One governance system across both streams. A single AI Transformation Office, a lean hub-and-spoke Centre of Excellence,[12] owns the master roadmap, runs the quarterly reviews, manages the portfolio, surfaces and consolidates shadow AI, and holds the authority to back, retire, and escalate that the VC model depends on.

4. Phased Implementation Roadmap

The model below is a template for putting the Hybrid Approach into practice. Proofs of concept (working prototypes built on real company data) recur throughout as the unit of validation that separates it from design-on-paper transformation. The timelines assume reasonable data and tooling readiness. Mid-cap organisations with legacy data architectures should add 2–3 months in Phase 1, with later phases shifting in step.

#	Phase	Bottom-Up Activities	Top-Down Activities	Key Outputs
1	**Foundation (Months 1–2)**	AI literacy training; no-code tooling; ideation workshop; first rapid PoCs	Stakeholder vision; domain scoping; benchmark analysis	Shared AI vocabulary; validated PoC backlog; draft transformation vision
2	**Discovery (Months 2–4)**	VC-style ideation sprints; prototyping on	Cluster use cases by domain; gap analysis; dependency mapping	Validated use case portfolio; de-risked AI

#	Phase	Bottom-Up Activities	Top-Down Activities	Key Outputs
		real data; feasibility scoring; stage gates		roadmap; prioritised investment plan
3	**Roadmap Alignment (Months 3–5)**	Business experts challenge redesigned process blueprints; champions identified	Process redesign blueprints; resource modelling; business case	Agreed hybrid roadmap; validated architecture; funded program
4	**Build & Scale (Months 5–18)**	Agile MVP delivery in ring-fenced environments; champions drive adoption; feedback loops; continuous PoC intake	Program governance; integration architecture; capability building	Live AI solutions; measurable KPI improvement; scaling playbook

5. Conditions for Success, and How Programs Fail

The Hybrid Approach is not self-executing. The conditions below are prerequisites, not refinements. Cut corners on any of them, and a program slides back into a single-mode failure pattern.

What must be in place. An executive sponsor with real authority over process, headcount, and investment. AI literacy paired with multi-framework regulatory readiness: Article 4 of the EU AI Act makes structured training a legal obligation for AI-facing staff, enforced by national authorities from 2 August 2026,[15] while GDPR and NIS2 cover the personal data and operational systems most use cases touch. Address one and ignore the others, and the exposure simply moves. Asset protection by design, so company IP cannot leak into public models. No-code tooling, so participants build without waiting on engineers. Real operational data, because synthetic data misleads on feasibility. VC-style governance with authority to back, retire, or escalate at each gate. And psychological safety, so employees can propose automating parts of their own roles. Without it, the portfolio avoids its highest-value opportunities.

How programs fail. They skip the literacy training, the most common shortcut, and now a regulatory exposure as well as a source of poor use cases and failed adoption. They treat AI as a technology project rather than a change in how work is done, which BCG names the defining error for CEOs who fail to capture value[7] and which shows up as the micro-productivity trap.[13] They treat governance as a one-time gate, and the portfolio goes stale within six months. They let the top-down stream override bottom-up input, which collapses the program back into pure top-down, and into the scaling failure seen in March 2026, when only 14% of enterprise AI-agent initiatives reached production scale.[14] Or they deploy no-code tools without real data, producing polished prototypes that reveal nothing about feasibility.

6. How LastingImpact.AI Works

LastingImpact.AI delivers the Hybrid Approach as two parallel streams joined by an integration gate, then a longer execution arc. Each part has explicit deliverables, governance, and value metrics.

The AI Leadership Program: top-down stream. We work with the C-suite and Board over three to four months to set the strategic, regulatory, and architectural scaffold. It is leadership-only, with no hands-on

prototyping, and runs five phases: diagnostic, AI literacy, operating-model reference architecture, governance and EU AI Act compliance, and the value-creation framework. **It gives the board control of the transformation** through a bounded Stage-1 blueprint: strategic mandate, governance V1, reference architecture, value-creation thesis, a compliance baseline spanning the EU AI Act (including the Article 4 board-oversight cadence), GDPR, and NIS2, an asset-protection regime, and a board-approved mandate to launch the bottom-up stream with named resourcing. It deliberately does not produce the prioritised use-case portfolio or the role-level operating model. Those depend on bottom-up validation and emerge from the integration gate.

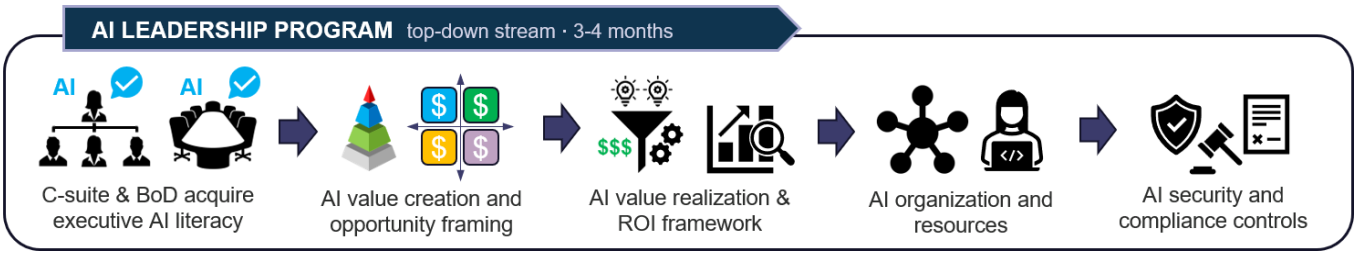

The AI Acceleration Program: bottom-up stream. Alongside the AI Leadership Program, or in sequence after it, we run a two-month Acceleration Program with cohorts of 10–20 domain experts. They move through an AI Academy (compliant with EU AI Act Article 4), a one-day ideation workshop, and a no-code sandbox where teams build and validate PoCs on real company data. A cohort typically produces 50+ structured use cases, 15+ validated prototypes, and one or two AI MVPs ready for production scoping.

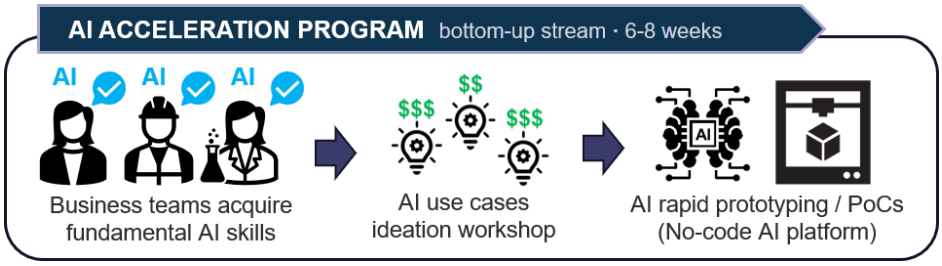

Hybrid Innovation Governance: the integration gate. An integration sprint after the first cohort, then a recurring quarterly review, maps validated use cases against the Stage-1 architecture under the VC-style discipline of Section 3. This is where **validated strategy becomes a deployable operating model**, captured in the Stage-2 blueprint: role-level redesign, the prioritised roadmap with validated feasibility, ROI projections grounded in PoC data, and the deployment plan. The Board approves both stages. Four workstreams then convert validated strategy into enterprise value: operating-model redesign;[4] [7] AI Governance V2.0 at regulatory grade for high-risk systems under EU AI Act Article 26; the Centre of Excellence built out as a permanent capability;[12] and quarterly value-realisation tracking so the roadmap never goes stale. Run together, the streams climb the outcome ladder the model is built for: first automation pilots, a mature data and AI execution capability, AI products at scale, an AI-first organisational redesign, and an advanced AI capability that compounds into competitive edge (see next page figure.)

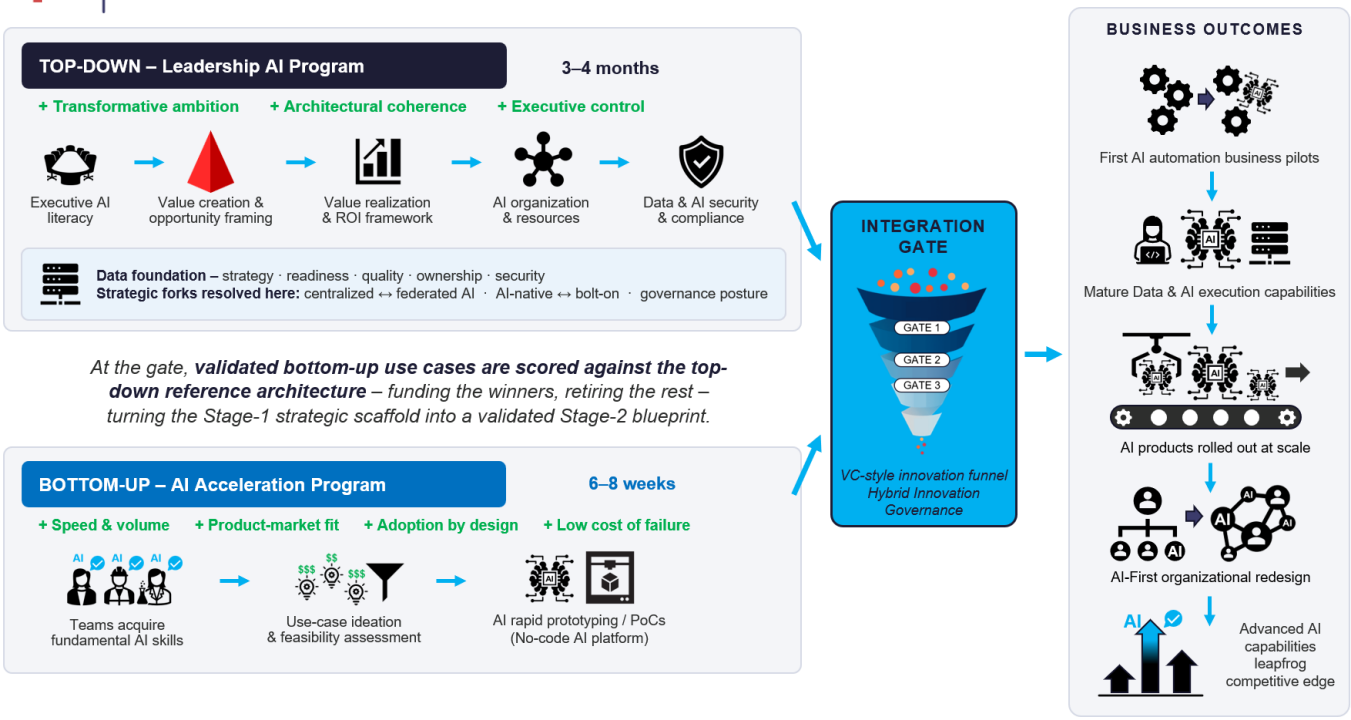

Proof from delivery. We have run this model inside live businesses and measured the results. Because these engagements are bottom-up, they show the firm delivers outcomes; they are not a claim that bottom-up alone suffices. In one, a European consultancy of 200-plus PhD-level bioengineers needed its AI innovation to come from its own domain experts rather than an outside "AI guru." Successive cohorts delivered two waves. The first lifted productivity on core process-configuration work by 50% across a large share of project volume.

The second conceived a new AI-native service line, viable only by fusing the firm's domain expertise with company-owned AI agents, and it opened a fresh revenue stream and a previously unreachable mid-market segment in pharma. In another, two real-estate-services companies under common ownership built a cross-functional workgroup that deployed shared AI under a federated data-governance model: predictive maintenance, a GenAI field-technician assistant, and joint demand forecasting. Unnecessary site visits fell 15–18%, first-time fix rates improved, and the workgroup became a durable in-house AI capability. The method also holds at global scale: under strict NDA, we have delivered it inside a pharma-equipment manufacturer of 16,000 employees and a pharma group of 40,000, with the same pattern of results. For a leader accountabl e for AI, all of it shows the same thing: trained domain teams turning into a capability the organisation keeps.

How we deliver, and what it means for you. Engagements are led by our founder, with AI transformation experts in ML, GenAI architecture, data governance, and regulatory compliance (GRC) brought in as each phase requires. Our approach is nimble with your resources: the AI Leadership Program asks five to six hours per month of each executive, the AI Acceleration Program two to three days of total effort per participant, and both calibrate to your operating calendar. Every organisation starts from a different place. Get in touch and we will work out the sequencing that fits yours.

About the author

Jérôme Zürcher is the Founder and Managing Director of LastingImpact.AI, an AI transformation firm that supports organisations through every stage of their AI maturity, from AI innovation acceleration to company-wide transformation. He is a lecturer in AI Strategy and the Fundamentals of AI and Machine Learning in executive MBA programs at ETH Zurich and Babson College, Massachusetts.

Sources

[1] Challapally, A., Pease, C., Raskar, R., & Chari, P. (July 2025). *The GenAI Divide: State of AI in Business 2025*. MIT Project NANDA.

[2] Bousquette, I. (April 18, 2025). *Johnson & Johnson Pivots Its AI Strategy*. Wall Street Journal.

[3] Yu, H., Greeven, M., et al. (August 2025). *Winning with AI: The Business Leader's Guide to AI from Strategy to Execution*. IMD Business School.

[4] Singla, A., Sukharevsky, A., Hall, B., Yee, L., & Chui, M. (November 5, 2025). *The State of AI in 2025: Agents, Innovation, and Transformation*. QuantumBlack, AI by McKinsey.

[5] Korst, J., Puntoni, S., & Tambe, S. (October 2025). *Accountable Acceleration: Gen AI Fast-Tracks Into the Enterprise*. Wharton Human-AI Research / GBK Collective, third annual study (n=800+ US enterprise leaders). ROI figures are self-reported by respondents.

[6] McKinsey & Company (February 2026). *The State of Organizations 2026*. Based on the McKinsey State of Organizations 2026 Survey, May–September 2025, n=10,018.

[7] Boston Consulting Group (January 15, 2026). *BCG AI Radar 2026: As AI Investments Surge, CEOs Take the Lead on Decision Making and Upskilling Themselves*. Based on a survey of 2,360 executives across 16 markets and 9 industries, including 640 CEOs.

[8] PwC (December 2025 / January 2026). *2026 AI Business Predictions*.

[9] Deloitte AI Institute (January 2026). *The State of AI in the Enterprise 2026: From Ambition to Activation — The Untapped Edge*. Survey of 3,235 leaders across 24 countries, August–September 2025.

[10] Korst, J. (April 2026). *Managers and Executives Disagree on AI — and It's Costing Companies*. Harvard Business Review.

[11] Lakhani, K. R., & Spataro, J. (March 2026). *The Last Mile Problem Slowing AI Transformation*. Harvard Business Review.

[12] Hub-and-spoke AI Centre of Excellence as the operating model most 2026 practitioner sources converge on. Named enterprise sources include Microsoft's AI CoE guidance and implementations by NTT DATA, Capgemini, and EY.

[13] Harvard Business Review (April 30, 2026). *How to Move from AI Experimentation to AI Transformation*. Introduces the term 'micro-productivity trap' for task-level AI gains that do not aggregate into firm-level value.

[14] March 2026 enterprise survey of 650 technology leaders on AI-agent pilot-to-production scaling, reporting 78% with pilots and only 14% at production scale; see Digital Applied, *AI Agent Scaling Gap*, March 2026, and related industry analyses including Gartner forecasts of 40% agentic-AI project failure by 2027.

[15] European Union (2024). *Regulation (EU) 2024/1689 (Artificial Intelligence Act), Article 4 (AI Literacy)*. In force since 2 February 2025; enforcement by national competent authorities and market-surveillance authorities begins 2 August 2026. See also European Commission AI Office, *AI Literacy Questions & Answers*.