COUNCIL ADVISORS

The Execution Chasm

Why AI Transformation Is an Operating System Problem

November 2026

SSA & COMPANY

TMG

HLG | HIGH LANTERN GROUP

Executive Summary

Most enterprises are losing the AI race. Not for lack of technology or budget: more than 85% of large enterprises already use AI[1], and spending keeps climbing. Only 5% of organizations report real financial gains from AI.[2] Buying the best frontier model or running another round of pilots won't close that gap.

The bottleneck is the corporate operating system — how a company senses a change, decides what to do about it, and improves.

The gap is not technical. It's an execution gap, and it's widening. BCG calls the leaders who've embedded AI into their operating models "future-built" and they're pulling away fast, while everyone else risks falling permanently behind as AI turns from experiment into infrastructure.[2] What's needed now is an integrated plan to close that gap, combining strategy, thoughtful operating-model design, and change management to move teams out of pilot purgatory and into P&L impact they can see.

Why It Matters Now: AI Is a Different Kind of Disruption

Most leaders keep treating AI like a faster version of the internet or an ERP rollout: buy the tool, train the users, wait for the productivity gain. AI is different — and that's why initiatives stall.

AI breaks that model in four ways.

- **AI hits individual tasks, not whole jobs or functions.** ERP and the internet re-engineered entire functions or connected people faster across a fixed structure. AI changes how a single task gets done, task by task, which creates friction no generic software license can fix and forces the organization itself to flex in real time.
- **The tools work only if humans and machines keep feeding each other data.** A one-time deployment doesn't build an advantage; a running feedback loop between people and machines does. That loop is what a competitor can't simply copy.
- **Every employee uses the tool differently.** Traditional IT assumed everyone would use the software the same way. AI use varies person to person,

so workflows built around one "correct" way of working break down immediately.

- **Differentiation must be built, not bought.** Importing a vendor's point solution can't create an advantage a competitor can also buy. The only lasting edge comes from redesigning your own processes around your own data.

The AI Paradox: Massive Investment, Elusive Returns

Global generative AI spending topped $37 billion in 2025, and enterprise conviction is high. AI spending is not the problem. The challenge is that the results aren't yet there.

- **The failure rate.** Roughly 80% of enterprise AI projects fail to deliver the business value they promised — twice the failure rate of conventional IT projects.[3]
- **Pilot purgatory.** Of every 33 AI proofs-of-concept a company launches, only four reach production. The other 29 stall or shut down, without ever recouping their cost.[4]
- **The widening divide.** "Future-built" leaders are posting 3.6x the total shareholder returns and 1.7x the revenue growth of laggards over three years.[2]
- **The private equity overhang.** The gap is sharpest in private equity, where roughly 29,000 portfolio companies are held globally, half of them for five years or more. Generic cost-cutting chatbots don't move that needle. Firms need real operational gains — to expand margins during the hold and prove to the next buyer that the AI-driven value survives the sale.

An Adaptive Challenge in an AI Costume

Most consultants diagnose this as a technology problem. It isn't. BCG's research puts the real split at 10% algorithm, 20% technology and data, and 70% people and process.[2] Most companies invert that, putting their effort into the 30% that is easy to buy (model subscriptions, software licenses, compute) while treating the 70% that decides the outcome as an afterthought.

The 10/20/70 Rule for Resource Distribution[2]

10%	20%	70%
algorithms	**technology & data**	**people & process**

Layering a powerful model onto a process nobody redesigned simply executes the same flawed workflow faster. It's the classic Lean failure, now with a chatbot attached. Technology produces little without changes to decision rights and to the management system underneath.

The fix starts with recognizing what kind of problem this is. Ronald Heifetz drew the line decades ago: technical problems get solved with existing expertise and the right tool. Adaptive challenges require people to change their values, beliefs, and behavior. By its nature, AI is adaptive. The old transformation model — unfreeze, optimize, refreeze — doesn't work here, because there's no stable state to refreeze into. Stabilization is temporary now, not permanent.

Additionally, the bottom knows before the top. AI disrupts work at the task level, so frontline employees run into its limits and its capabilities well before that reality shows up on a leadership dashboard, potentially creating a real blind spot at the top.

The Strategic Weapon: Accelerating Learning Velocity

If frontier models are becoming a commodity anyone can rent, where's the advantage? It's Learning Velocity: how fast an organization can cycle through gathering data, drawing insight, deciding, and acting.

Today, the real bottleneck is the Latency Gap: the time between an organization sensing a shift and doing something about it. A model can flag a supply chain anomaly in milliseconds. If it then takes three weeks of meetings to redirect inventory, the model's speed is wasted.

Organizations experienced in business process transformation redesign corporate structures to close that gap. The redesign moves organizations from command-and-control hierarchies to product-centric, outcome-based models where the people closest to the AI signal also have the authority to act on it.

Bridging the Chasm: The Six Pillars of AI Operating Model Transformation

Point solutions fail because they treat a multi-dimensional problem as a single fix. A strategy firm hands over a roadmap with no way to execute it. A technology vendor ships a tool the business won't adopt. SSA, part of Council Advisors, has an end-to-end model that closes both gaps by putting strategy, operating-model design, and change management under one accountable team.

01 **Strategic Clarity.** Define quantified business outcomes at the enterprise level, then align value pools to the P&L before investing capital or writing code.

02 **Operating Model Design.** Completely restructure workflows, roles, and governance around AI. Compress the Latency Gap by relocating decision authority to match technical velocity.

03 **Tech & Data Foundations.** Remediate data quality — which 70% of enterprises cite as their top AI challenge — to ensure models run on clean and dependable institutional data.[5]

04 **Talent & Change Enablement.** Close the global skills gap through parallel role redesign and upskilling, managing the human aspect of the transition as an adaptive challenge that requires shifting behaviors rather than just tool training.

05 **Governance & Risk Frameworks.** Build visible oversight mechanisms, alignment incentives, and ethical guardrails into the operating model from day one. (Only 1 in 5 companies has a mature model for autonomous agents.)

06 **Scale & Sustain.** Treat deployment as an ongoing operational cycle rather than a one-time IT launch. Implement stage-gating and continuous improvement cadences (PDCA — Plan/Do/Check/Act) to ensure the system permanently learns.

These aren't new capabilities invented for the AI moment. They're the same Lean, Six Sigma, and management-system disciplines that can be applied to smart business transformation.

Garrett Lord on Why AI Execution Fails — and What Winners Do Instead



Garrett Lord is the co-founder and CEO of Handshake. In January 2025, Handshake launched Handshake AI, which turned its network of PhDs and domain experts into expert evaluation and training data for the frontier AI labs. Sixteen months later, it works with eight leading labs and has passed $1 billion in gross annualized revenue, one of the fastest-growing businesses in AI. Handshake now brings that same expert rigor to enterprises: the evaluations and training environments that measure whether AI can do real work, and the models trained on them, built on each company's own expertise and data. Garrett was named to the TIME100 AI in 2026.

Few operators see both sides of that market. In our Q&A with Garrett, we talked with him about why so much corporate AI investment fails to pay off, and what the companies pulling ahead are doing differently.

Executives keep telling you investment in AI isn't delivering returns. Why?

Because they're managing it like software when it's nothing like software. The instinct is to stand up an agent, put a few engineers on it, and debug the problems as they come up. But an intelligent system isn't code. It gives you a different answer from one run to the next, and it shifts every time the model or the context changes. This means that small failures don't happen in isolation. They compound and cascade through everything downstream. Oftentimes this can be hard to measure, and if you can't measure it, you can't improve it — you're just spending more to fail in ways you can't see yet.

So, what are the companies getting real value doing differently?

They stop debugging and start improving. Think about how Waymo works. It doesn't fix self-driving one bug at a time. It locks a version, distributes it, measures every mile, and improves the next release. That's a control loop, and it starts with a clear definition of what "good" actually looks like. That is the outcome of agreement among your best

experts, measured against the real distribution of work coming in the door. Get a handful of experts to agree on the exact outcome and the best responses, and now you have a target you can measure against.

Take an accounting or legal practice, where the entire product is expert judgment applied to documents. The instinct is to use AI to produce those documents faster. Speed can be a helpful by-product, but the real opportunity is to teach a system to achieve the result your best people would reach, every time.

The labs have effectively unlimited computing power, yet they still rely on human experts. What does that tell you?

It tells you where the value is. If the most sophisticated AI companies on earth conclude that expert human judgment is the scarce input to a better model, that should tell an ordinary company something about where its own investment should go. When someone asks how they'll know their agents are working, the weak answer is "someone on the team checks it." The real answer is experts, people who can read a long, complex output, bring real domain knowledge, and tell you exactly where the reasoning breaks. Humans in the loop need to have real expertise.

> ***Your data is the whole game, and most companies are sitting on it without doing anything about it.***

If expert judgment is the scarce input, where should a company be spending and who owns this?

First, understand this is not an IT project — framing it as one is the clearest reason it stalls. An integration has an end. You scope it, you deliver it, the team signs off and moves on. An intelligent system is never finished. It needs someone measuring it and improving it every month, and if you've framed the work as delivery, nobody is left holding that once it's live.

The shift is from automating the task to capturing the judgment of the person. That doesn't mean replacing them. It means capturing how your

best people actually make decisions and keeping a human in the middle of the loop rather than at the end. Someone in the business has to own it, because the honest way to measure it is in margin, not in tickets closed. The workflow already costs you something today; this is about doing it better, on your own terms.

Where does this go? What does winning look like three years out?

The companies that win won't be clients of AI labs — they'll behave like labs themselves. A lab is really only three things: compute, which is easy to get; algorithms, which are becoming available to everyone; and data, which only you have. Two of those are commodities. Your data is the whole game, and most companies are sitting on it without doing anything about it.

Behaving like a lab means running a fixed cycle rather than making a bet. You don't pick the best model and hope it holds. You build a deployment cycle around your own data, compare what's available, and keep your quality and cost advantage as the models underneath keep changing. The winners are the ones who start turning their own data into a measurable advantage while everyone else is still evaluating models.

Key takeaways — where to start:

- Start on your own data; it's the only part that your competitor can't buy.
- Define "good" first. Your best experts, real examples, one agreed outcome.
- Build the loop around it. Lock a version, measure it against real work, improve the next release.
- Budget for expert judgment, not just tools.
- Give it to a business owner who already carries the cost of that workflow and measure it in margin.

Conclusion: The Critical Question for This Quarter

The window for AI transformation is narrowing. Fast-forward two years, and the gap between "future-built" leaders and everyone else will be extremely hard to close. The conversation worth having this quarter isn't which model to standardize on or vendor to sign. It's whether your operating system can absorb the technology, relocate decision rights, and keep learning.

Designing transformation from the inside out is the answer, addressing process, decision rights, management systems, and continuous improvement in an integrated way. Done well, that will be the difference between AI ambition and AI impact.

Authors



Scott Miller
Chairman
Council Advisors



David Niles
CEO
Council Advisors

References

1. McKinsey, *The State of AI 2025: Agents, Innovation, and Transformation* (November 2025)
2. BCG, *The Widening AI Value Gap: Build for the Future* (September 2025)
3. RAND Corporation, *The Root Causes of Failure for Artificial Intelligence Projects and How They Can Succeed* (2024)
4. Lenovo with IDC, *CIO Playbook 2025: It's Time for AI-nomics* (February 2025)
5. McKinsey, *A data leader's technical guide to scaling gen AI* (July 2024)

Related Council Advisors Reading

Change Management in the AI Era — Part 2: Adaptive Leadership in Practice (May 2026)

Structure Is Strategy: How Organizational Design Determines Learning Speed (April 2026)

Learning Velocity: The Competitive Advantage Organizations Should Not Ignore (March 2026)

Change Management in the AI Era — Part 1: Leading Change When the Future Won't Stand Still (March 2026)

Leadership Principles: Closing the AI Execution Gap in PE (March 2026)



COUNCIL ADVISORS

1.212.332.3790
www.counciladvisors.com





