
The Transform Imperative

Why the Real Value of AI Lies Beyond Efficiency, and What C-Suite Leaders Must Do About It

Gerald C. Kane, Ph.D.

C. Herman and Mary Virginia Terry Chair in Business Administration
Department Head, Management Information Systems
Terry College of Business, University of Georgia

© 2026 Gerald C. Kane. Licensed under CC BY-NC 4.0: share and adapt for noncommercial use with credit to the author; do not present as your own. Full terms on the last page.

Executive Summary

Most enterprises are leaving the majority of AI's value on the table — and they do not yet know it.

McKinsey's 2025 *State of AI* report found that 88% of organizations now use AI in at least one function. Only about 6% of organizations report it contributing more than 5% to EBIT. That gap — between adoption and value — has become the defining strategic problem of the AI era.

The gap is not a technology gap. The technology is commoditizing. The gap is a **framing gap**. Most enterprises are using AI to do existing work faster and cheaper. That captures real but bounded value, and competitors can replicate the move within a year. The enterprises pulling decisively ahead are using AI to do something categorically different: to change what the organization does, not just how fast it does it. They are redesigning workflows, reconceiving products, and reshaping business models around capabilities that did not exist before.

Behind the framing gap sits a deeper problem this paper builds on. General-purpose technologies do not produce productivity gains immediately on deployment. The gains arrive later, after complementary investments — in workflow redesign, organizational restructuring, workforce reskilling, new product co-invention — mature enough to use the technology fully. Brynjolfsson, Rock, and Syverson named this the productivity J-curve: in the early years measured productivity dips while intangible complements accumulate; later, when the complements mature, productivity rises faster than it should. AI is in the dip now. The strategic question for the C-suite is not whether the rise is coming. It is whether the organization can fund the intangible middle long enough to reach it — and whether the executive team has the language to defend the dip to boards and analysts who do not yet understand why the curve looks the way it does. Most do not. That is the deepest reason most enterprises will not capture transform-layer value: not because they picked the wrong strategy, but because they could not hold the position long enough for the strategy to mature.

This paper argues that the distinction between these two postures is the most important strategic choice the C-suite faces in the next twenty-four months. It uses a framework introduced by Harvard scholar Shoshana Zuboff in her 1988 book *In the Age of the Smart Machine* — and extended by subsequent academic and practitioner research — to name the three capacities of any information technology:

- **Automate** — using technology to replicate existing work with less human labor.
- **Informate** — using technology to generate new visibility, information, and learning about the work itself.
- **Transform** — using technology to enable fundamentally new ways of working, new products, and new business models.

Transform is where durable advantage lives. But transform is not a skip-ahead. It builds on the other two. Organizations that attempt to transform before they have learned to automate and informate produce demo-ware that never leaves the pilot. Organizations that stop at automate produce cost reductions competitors will replicate within a year. The work is sequenced: automate

first, informate to make the automation worth something, and then transform the business on the foundation the first two created.

This framing has two implications the C-suite cannot responsibly delegate.

First, AI strategy is a portfolio allocation decision. Most enterprises today are heavily overweight on automate projects, lightly weighted on informate projects, and have no transform projects at all. Rebalancing the portfolio toward transform is a capital-allocation and organizational-design decision that cannot be handed to IT or to a functional leader. It belongs to the executive team.

Second, transform-level decisions require personal fluency. BCG's January 2026 AI Radar found that only about 15% of CEOs — the "Trailblazers" — are spending more than eight hours a week on their own AI upskilling. Those leaders' companies are pulling decisively ahead. Transform requires an executive who can tell, from direct experience, what AI can and cannot yet do. That judgment cannot be outsourced to a consultant, a vendor, or a functional leader — because none of them share the accountability for the enterprise as a whole.

The paper is organized into three parts. **Part I** establishes the paradox of widespread adoption and narrow value capture, and introduces the automate / informate / transform framework. **Part II** walks each of the three layers in turn, with named cases of organizations succeeding and failing at each — from Klarna's automate-layer reversal to Moody's informate-layer redesign of credit analysis to John Deere's transform-layer shift from equipment manufacturer to data and subscription platform. **Part III** describes what the C-suite must personally do to move the enterprise up the stack — including a portfolio allocation model, a fluency-building regimen, and a governance structure that keeps the transform agenda from being starved by automate urgency.

The conclusion is straightforward. The AI era will not reward enterprises that use the technology to do less of what they already do. It will reward enterprises that use it to do what they could not do before. Getting there requires a C-suite that personally understands the technology well enough to tell the difference.

Part I — Framework

I. The AI Paradox: Massive Adoption, Narrow Value Capture

The data on enterprise AI, taken together, describes something most C-suites have been reluctant to say out loud: we are spending record amounts on a technology most of us have not yet figured out how to monetize.

McKinsey's November 2025 *State of AI* survey found that 88% of organizations now use AI in at least one function, up from 78% a year earlier, 72% the year before that, and 55% three editions back. But only about one-third have scaled AI across the enterprise. The other two-thirds remain stuck in pilots. Just 39% of organizations report *any* EBIT contribution from AI, and most of those say it accounts for less than 5%. Roughly 6% of organizations — the group McKinsey calls "AI high performers" — report more than 5% of EBIT attributable to AI. The gap between the general population and the high performers is not marginal. It is the entire ball game.

Agentic AI, the newer and more capable generation of AI systems, tells the same story at an earlier stage. In the same survey, 62% of organizations said they are experimenting with AI agents. Only about 23% are scaling them. And BCG's *Widening AI Value Gap* survey of 1,250 senior executives reached a starker version of the same conclusion: 60% of companies are capturing hardly any material value from AI, reporting minimal revenue and cost gains despite substantial investment.

The paradox sharpens when you look at what has been happening in the labor market. PwC's 2026 *Global AI Jobs Barometer*, based on more than a billion job postings across 27 countries, found that companies in the most AI-exposed sectors recorded 34% productivity growth in 2025 relative to 2018, against 24% for the sectors least able to use AI. Job postings requiring AI skills grew 69%, roughly eight times the 9% growth of the overall jobs market. Workers with AI skills now command a 62% wage premium, up from 57% a year earlier. The technology is clearly working somewhere.

So we have an asymmetric outcome: the technology is producing massive productivity gains at the industry level, while most individual organizations deploying it report little to no EBIT impact. A small number of firms are capturing almost all the value. Everyone else is running pilots.

Why this is happening

The instinct in most boardrooms is that the gap must be technical. The high performers must have better models, better data, better engineers. If we can catch up on the inputs, we will catch up on the outputs.

The evidence says otherwise. The models are available to everyone — the same foundation models from OpenAI, Anthropic, Google, and Meta are accessible through any enterprise contract. The data infrastructure patterns are well-documented. The engineering talent is tight but not unobtainable. What differentiates the high performers is not their stack. It is what they are asking the stack to do.

McKinsey's own analysis of its high-performer group makes this point directly: the strongest single correlation with EBIT impact from AI is **fundamental workflow redesign**. High performers are

roughly 2.8 times more likely than their peers (55% vs. 20%) to report redesigning end-to-end processes. Senior-leadership ownership matters almost as much; high performers are three times more likely to say their top executives actively champion AI and personally role-model its use. The differentiators are organizational and managerial, not technical.

BCG reached the same conclusion by a different route. After studying hundreds of AI transformations across its client base, the firm codified what it calls the "10-20-70 rule": 10% of the effort required to capture value from AI belongs to the algorithms, 20% to the technology and data infrastructure, and 70% to the people and process changes. Organizations that invert that ratio — spending 70% of their effort on algorithms and infrastructure, and 10% on people and process — produce pilots that technically work and commercially fail.

A pattern with deep historical roots

There is nothing new about this pattern. Every general-purpose technology in modern economic history has produced the same early gap between adoption and measurable value. Stanford economist Paul David's 1990 paper *The Dynamo and the Computer* traced it carefully for electricity. The lightbulb was invented in 1879. By 1900, electric motors accounted for less than 5% of factory mechanical drive. Productivity gains from electrification did not show up in the statistics in any meaningful way until the 1920s — roughly four decades after Edison's first central generating station opened in 1882.

David's explanation is the one this paper will return to throughout. The early adopters of electricity did not actually change what their factories did. They replaced the steam engine at the center of the factory with a large electric motor, kept the mechanical power distribution system of belts and pulleys, and kept the factory layout that the belts and pulleys had required. The technology had changed; the organization had not. It took another generation of factory owners — and another generation of factory architects and electrical engineers — to design the factory that electricity actually made possible: a single-story, modular plant with individual motors on each machine, laid out according to the flow of materials rather than the geometry of power distribution. That redesign is where the productivity gains lived. Not in the motor.

Erik Brynjolfsson, Daniel Rock, and Chad Syverson formalized this pattern in a 2021 paper in the *American Economic Journal: Macroeconomics*. They call it the "Productivity J-curve." General-purpose technologies like AI require substantial complementary investments — in business process redesign, co-invention of new products and business models, workforce reskilling, and organizational restructuring — that are mostly intangible and not captured in conventional productivity measures. In the early years, productivity appears to dip: the organization is investing in complements, and the payoffs have not yet arrived. Later, when the complements are in place, productivity grows faster than it should. The curve dips, then rises. Organizations that cannot fund the intangible middle — or that lose patience before they reach the rise — never see the returns.

This is the frame this paper is built on. The productivity gains from AI will not come primarily from the algorithms. They will come from the complementary investments the algorithms enable. And the C-suite's most consequential job in the next twenty-four months is to make sure those complementary investments actually happen.

II. Three Orders of Technology: Automate, Informate, Transform

The most useful framework for making sense of the enterprise AI paradox was written thirty-eight years ago, by someone who had never seen a large language model. In 1988, Harvard Business School professor Shoshana Zuboff published *In the Age of the Smart Machine: The Future of Work and Power*, a book built on a five-year ethnographic study of eight organizations adopting information technology — including a pulp and paper mill she called Piney Wood, where she watched operators move from a physical, hands-on relationship with the production process to a new one mediated entirely by screens and data.

Zuboff's central observation was that information technology has a dual capacity that earlier generations of mechanization did not. Mechanical automation substitutes for human labor: a loom replaces a weaver, a steam engine replaces a water wheel. The logic is one-directional — the machine does what the human used to do, faster and cheaper. Information technology does that, too. But it does something else simultaneously. **It produces information about the very processes it automates.** The sensors that control a paper machine also generate a continuous record of temperature, pressure, moisture, throughput, and variance. The terminal that replaces a bank teller also captures every keystroke, every transaction, every pattern of customer behavior. The cash register that automates checkout also produces the data stream that makes inventory management, dynamic pricing, and demand forecasting possible.

Zuboff's term for this second capacity was **informating**. Information technology, she argued, does not just automate; it informs. It "translates activities, objects, and events into information," rendering them "visible, knowable, and shareable in a new way." It creates what she called an "electronic text" — a continuously accruing record of what the organization does, accessible to anyone with the right access to read it.

The argument of her book was that managers faced a choice. They could deploy information technology to automate — using the data to monitor, control, and deskill the workforce, treating information as an instrument of centralized supervision. Or they could deploy it to informate — using the data to push decision-making authority down to the people doing the work, treating information as the raw material of broader organizational intelligence. The technology was the same. The strategic posture was different.

From two orders to three

Zuboff's framework was descriptive of the computing technology of the late 1980s. It did not explicitly contemplate a third order — a capacity for technology to enable not just the automation or informating of existing work, but the fundamental reconception of what the organization does. That third order was added by subsequent academic research, most directly by Bruce Dehning, Vernon Richardson, and Robert Zmud, whose 2003 paper in *MIS Quarterly* — "The Value Relevance of Announcements of Transformational Information Technology Investments" — explicitly built on Zuboff to define a third strategic role for IT alongside automate and informate: **transformational IT**.

The Dehning paper did something important beyond naming the category. It measured the effect. Using an event-study methodology on stock-market announcements of IT investments, the authors found that firms announcing *transformational* IT investments — investments that changed what the

firm did, not just how fast it did existing things — experienced positive abnormal stock returns. Firms announcing purely automating IT investments did not. The market, at least in the aggregate, correctly distinguished investments that would reshape competitive position from investments that would merely reduce cost. Cost reductions were already priced in. Reshaping was not.

A parallel line of argument reached the same conclusion through a different lens. N. Venkatraman, writing in *Sloan Management Review* in 1994, proposed a five-level hierarchy of IT-enabled business transformation, ascending from "localized exploitation" (using IT to improve individual functional tasks) through "internal integration" and "business process redesign" to "business network redesign" and finally "business scope redefinition." Each level required more organizational change and produced more durable advantage. Venkatraman's core finding was that most organizations stopped at the first two levels — improving individual tasks and integrating them — and never climbed to the levels where competitive position actually changed.

More recently, Harvard Business School professors Marco Iansiti and Karim Lakhani brought the same argument into the AI era. Their 2020 book *Competing in the Age of AI* describes what they call the "AI factory" — a new operating architecture in which algorithms, data pipelines, and software sit at the center of the firm and remove the traditional constraints on scale, scope, and learning that have limited business growth for centuries. Iansiti and Lakhani's claim is not that AI will help traditional firms do their existing work better. It is that AI enables a different kind of firm altogether — one that can serve, in their Ant Financial example, more than ten times the customers of the largest U.S. bank with less than one-tenth the headcount. Ant Financial is not a faster bank. It is a structurally different organization.

The three orders, defined for AI

The automate / informate / transform framework, updated for the AI era, produces three distinct categories of AI deployment. Each has a different logic, a different ROI profile, a different risk posture, and — crucially — a different kind of leadership requirement. The relationships among them are stacked, not sequential: each higher layer rests on the one below it.

The framework can also be read across its operational dimensions. The table below summarizes how the three orders differ on the questions an executive team will actually ask when deciding where to invest.

Dimension	Automate	Informat	Transform
Core question	How do we do this faster?	What can we now see that we couldn't see before?	What could we do that we couldn't do before?
What changes	Cost structure of existing work	Decision quality and organizational learning	Products, workflows, business model
Time to value	3–9 months	6–18 months	18–36 months
Dominant risk	Quality regression; brand damage	Decision paralysis; data quality exposure	Strategic misfire; cultural rejection
Competitive durability	Low — replicable within 12 months	Medium — depends on data and culture	High — compounds over years
C-suite role	Approve budget; monitor ROI	Set governance; model use	Own the thesis personally

The three orders are not mutually exclusive, and they are not competing strategies. They are a stack. Every enterprise deploying AI will do some of each. The question — the question this paper is really about — is what the *mix* should be, how the three layers reinforce each other, and which of them the C-suite personally needs to own.

III. The Layered Logic: Why Transform Builds on the Other Two

A natural reading of the previous section would be: transform is where the money is, so let's go straight there. That reading is wrong. It is the single most common strategic error in enterprise AI today.

Transform is not a standalone move. It is the top layer of a stack that rests on the other two. An organization that has not learned to automate competently cannot sustain a transform initiative, because the underlying workflows are too fragile. An organization that has not built informate capacity — the data infrastructure, the analytical muscle, the management habit of actually using what AI makes visible — cannot tell whether its transform bets are working. The dependency is structural rather than temporal. Each layer needs the substrate of the one below to be functional before the one above can stabilize, but mature organizations build the three in parallel — John Deere did not finish its automate program before deploying sensors and analytics, and Moderna integrated the three from founding. What the framework requires is that no transform initiative be launched without the informate substrate to interpret it and the automate competence to absorb the workflow changes — not that the layers be completed in strict order.

The lesson of the steam-powered electric factory

The clearest illustration of why the sequencing matters comes from the history Paul David described. In the first generation of factory electrification, in the 1890s, owners replaced the steam engine at the center of the factory with an electric motor. Everything else stayed the same: the central power source, the mechanical distribution system, the factory layout, the organization of labor around the mechanical constraints. The electric motor was more efficient than the steam engine on a per-horsepower basis, but the gains were modest — single digits. Most factories did not bother, which is why electrification looked like a failure for thirty years.

What changed in the 1920s was not the motor. It was the factory. A new generation of owners — and a new generation of factory architects and electrical engineers — designed plants around what electricity actually made possible: individual motors on each machine, single-story layouts driven by material flow rather than power distribution, new organizations of labor, new product categories that could only be manufactured in the redesigned facilities. Productivity in American manufacturing grew faster in the 1920s than in any decade before or since. Electrification accounted for roughly half of it.

The point is not that the early adopters were stupid. They were doing what most enterprises do with a new general-purpose technology: layering it on top of the organization that had been built for the old technology. They were *automating*. They were not yet informing or transforming, because they had not yet built the complementary capabilities — the design, the data, the labor model, the culture — that would make the other two layers possible. It took a generation of automate-layer learning before the transform-layer imagination could actually be executed.

The same pattern is playing out with AI, compressed

The AI parallel is almost exact. Most enterprises in 2026 are in the "steam-powered electric factory" phase. They have bought enterprise licenses to a foundation model. They have deployed it into customer service, code completion, document summarization, expense reporting. The old workflows are still in place; the old org charts are still in place; the old incentive structures are still in place. The AI has been dropped into the middle of all of it as a faster version of existing tasks. The productivity gains are real. They are not transformative.

The enterprises that are pulling ahead look different. They have not skipped the automate phase. They have gone through it and built the informate and transform capacities on top of it. They have used the automate phase to generate the data, the organizational muscle, the managerial habits, and — critically — the personal executive fluency that makes the next two layers possible.

The Sequencing Principle Automate produces the data and the organizational muscle. Informate produces the visibility and the managerial judgment. Transform produces the new products, workflows, and business models. Organizations that try to skip directly to transform produce demo-ware. Organizations that stop at automate produce commoditized cost reductions. The prize belongs to organizations that work the full stack.

The compression rests on three structural conditions that did not hold for prior general-purpose technologies. First, the digital substrate that AI requires is mostly already built. Cloud computing, enterprise data warehouses, broadband, mobile workforces, the systems-of-record that AI needs to read from — these took roughly twenty-five years to construct between 2000 and 2025. Electricity had to wait for factory architects, electrical engineers, and a generation of new construction; AI inherits an industrial economy that is already digitized. The substrate work that took the previous generation a decade is the starting position for this one.

Second, the managerial muscle for digital change is now in its third or fourth iteration. The leaders running enterprises in 2026 have lived through the web, mobile, cloud, and social transformations. The change-management vocabulary, the governance patterns, the workforce-transition playbooks do not have to be invented; they have to be applied. That is faster work than the work the computer adopters of the 1990s were doing, and dramatically faster than the work the factory owners of the 1900s were doing.

Third, the capability frontier of AI itself is moving on a cycle prior general-purpose technologies did not run on. Frontier model capabilities that did not exist eighteen months ago are now table stakes. Firms that wait for the technology to stabilize before beginning the complementary work will find that the technology has stabilized at a different level than the one their planning was built around. The compression is partly a feature of the inherited substrate. It is also partly a forced strategic deadline: even firms that could afford a slower timeline cannot afford to give competitors who move faster a structural lead that will compound over the same window.

These conditions are real, but they bound the compression rather than guarantee it. The substrate is in place; the human and organizational adaptation is not. The firms that capture the compressed timeline are the ones whose executive teams can absorb the change at the speed the substrate now permits, and the firms that fail will fail because the human side of the J-curve cannot be compressed the same way. The forecast in this paper is that productive enterprises will see the productivity rise within twenty-four months. It is not that every enterprise will be able to participate.

The intangible middle

Brynjolfsson, Rock, and Syverson's Productivity J-curve describes exactly this sequencing in economic terms. In the early phase of a general-purpose technology, organizations are investing heavily in intangible complements — process redesign, new organizational structures, workforce reskilling, co-invention of new products and business models — that do not show up as capitalized assets in the national accounts. Measured productivity dips because the inputs are real but the outputs have not yet arrived. Later, when the intangible investments mature, measured productivity overshoots. The curve is not linear. It is a J.

The practical implication for a C-suite is that there will be a period — in most cases one to two years — during which AI investment looks like it is underperforming. Finance will notice. The board will notice. Analysts will notice. That is the period during which the intangible complements are being built. Cutting investment at that point, or pulling the plug on the transform portfolio to chase automate-layer cost savings, locks the organization into the bottom of the J. It is the single most common failure mode in AI transformation, and it is almost always justified as prudence.

Recognizing the J-curve is the beginning of the C-suite's defensive job. Holding the organization on the curve long enough to reach the rise is the rest of it.

Why this paper is about the C-suite

The J-curve framing has an implication that turns the rest of this paper into an argument about leadership rather than technology. The intangible middle of the curve — the period during which complements are being built but the productivity gains have not yet arrived — is unavoidable. Every general-purpose technology has produced one. The question is not whether the dip happens. The question is whether the executive team is positioned to hold the organization on the curve through the dip.

That position requires three things the rest of this paper will return to repeatedly. First, a portfolio of investments diversified across the three layers, so that automate-layer savings fund the informate and transform layers rather than being returned to the P&L. Second, a governance structure that does not review transform initiatives on the same cadence as automate initiatives, because automate metrics will always look better quarter to quarter than transform metrics will look across multiple years. Third, an executive team with enough personal fluency in the technology to recognize, defend, and explain the dip when it appears — not as a failure to be corrected but as the structural feature of every general-purpose-technology adoption that has come before.

The single most common way enterprises fail at AI is not a strategic mistake. It is a discipline failure. The strategy is articulated; the early investments are made; the complementary work begins; the curve dips; finance notices; the board notices; analysts notice; and the executive team — without the language to defend the dip — pulls the plug on the long-cycle bets to chase the short-cycle savings. The transform layer dies, the automate layer continues, and eighteen months later the company is structurally indistinguishable from its competitors. That is the failure mode the rest of the paper is written to prevent.

Part II — The Three Layers in Practice

IV. The Automate Layer: Necessary, Visible, Over-Indexed

Automate is the layer most enterprises reach for first, and for good reason. It is concrete. It is measurable. The ROI calculation is straightforward — hours saved, calls deflected, headcount avoided, cycle times reduced. The business case writes itself, and the case can usually be defended in a quarterly operating review.

For a sequence of good reasons, most of the visible enterprise AI deployments of the last two years have lived here. Customer service chatbots handling tier-one inquiries. Code-completion tools speeding up developer throughput. Document summarization in legal and compliance. Automated expense report classification. Invoice processing. Contract review. These are real, useful applications. They produce measurable cost reductions. They are the right first move for most organizations.

The question is not whether to do automate work. It is whether to stop there.

Where automate delivers, and how much

The data on automate-layer AI is clearer than the data on any other layer, because automate projects are easy to measure. Customer service is the most mature use case: vendors and internal deployments consistently report call automation rates of 40-90%, cost reductions of up to 50%, and handling time reductions from 9 minutes to 2-7 minutes. Document processing has shown 95%+ accuracy rates on well-structured workflows and 6-12 month payback on typical deployments. Code-completion tools routinely produce 20-55% productivity improvements for individual engineers.

These are real gains. They fund the rest of the program. An organization that cannot execute automate-layer AI competently will not have the credibility, the budget, or the organizational confidence to pursue the layers above.

But these gains have two properties the C-suite needs to understand clearly. First, they are bounded. Automate-layer value is capped at roughly the cost of the work being automated, minus the cost of running the AI. If a function costs \$10 million a year and AI reduces that to \$6 million, the annual value is \$4 million. That is a real \$4 million; it is not a \$40 million number, and it will not compound. Second, these gains are replicable. The same foundation models that let your organization cut 40% off the customer service cost structure are available to every competitor. Within twelve to eighteen months, the cost curve rebases for the entire industry, and the competitive advantage is gone.

The Klarna reversal, and what it actually teaches

The most instructive automate-layer case of the last two years is Klarna's. In February 2024, Klarna announced that an OpenAI-powered customer service assistant was handling 2.3 million customer conversations in its first month — equivalent to the workload of 700 full-time agents — with response times dropping from 11 minutes to under two and a projected \$40 million profit

improvement. The company instituted a hiring freeze and positioned itself as the leading enterprise example of AI-first customer service. Headcount fell from roughly 5,500 at the end of 2022 to about 3,400 by late 2024, a 38% reduction.

By mid-2025, the strategy had publicly unwound. CEO Sebastian Siemiatkowski admitted that "cost unfortunately seems to have been a too predominant evaluation factor" and that "what you end up having is lower quality." Klarna began rehiring customer service staff, in many cases as flexible contract workers. Gartner had separately predicted, in a February 2026 release built on a survey of customer service and support leaders, that "by 2027, 50% of companies that attributed headcount reduction to AI will rehire staff to perform similar functions, but under different job titles."

The lesson most commentators drew was that Klarna "replaced humans with AI too soon." That is true as far as it goes, but it misses the more useful point. Klarna's mistake was not deploying AI to automate customer service. The mistake was **treating the automate-layer deployment as if it were the whole strategy**. The company optimized aggressively for the measurable variables — cost per interaction, resolution time, deflection rate — and lost track of the variables that were harder to measure but strategically decisive: empathy, de-escalation, brand trust, the quality of the customer relationship that makes a fintech a fintech rather than a commodity payment processor.

Read through the Zuboff framework, Klarna did automate-layer work, produced real automate-layer gains, and then treated those gains as if they were transform-layer gains. The actual transform move in customer service — the reconception of what customer service is *for* in an AI-enabled world — requires thinking about the combination of AI and human that produces a better customer relationship than either alone, not about the substitution of one for the other. Klarna is now doing that work, belatedly, with a smaller team and a chastened leadership.

When automate is the right move

None of this argues against automate-layer AI. It argues for *naming it correctly*. Automate is the right strategy when three conditions are met: the work to be automated is well-understood and stable; the quality floor can be held by human oversight in the cases the AI cannot handle; and the savings are being reinvested into the informate and transform layers, not returned wholesale to the P&L. Organizations that satisfy those three conditions will capture the automate-layer gains and build the foundation for the layers above. Organizations that satisfy only the first two will look like winners for a year and then watch their position erode.

The C-Suite Test for an Automate Project • Is the work being automated stable and well-understood enough that AI can handle the common cases with documented quality? • Is there a human-in-the-loop design for the cases AI cannot handle, and is it staffed? • Are the savings being reinvested into informate or transform initiatives, or being swept to the P&L? • Can you answer, in your own words, what the automate project does *not* do that matters strategically? If the answer to the last question is "nothing," the project is probably automating something that did not need to be done in the first place — and the informate and transform conversations are still ahead.

V. The Informate Layer: The Underappreciated Middle

Informate is the layer the current AI discourse has largely skipped. The automate conversation is everywhere. The transform conversation is fashionable. The informate conversation — the quieter, more foundational work of using AI to generate new visibility, new information, and new

organizational learning — is routinely left out of board presentations, which is part of why so many transform initiatives fail.

This is a mistake. Informate is where the organization learns what to transform.

Zuboff's original insight — that information technology generates information about the processes it automates — is more powerful with AI than with any previous generation of technology. A customer service AI that handles tier-one inquiries is simultaneously generating a structured, searchable, analyzable record of every customer interaction, every question asked, every resolution offered, every escalation triggered. A coding assistant is generating a record of every pattern engineers struggle with, every workaround, every place the codebase resists modification. A document-processing AI is surfacing, for the first time at scale, the structure of work that had lived in individual heads.

This is the electronic text Zuboff described, now generated at a volume and a resolution that was not possible in 1988 and was still not possible in 2018. The question is whether the organization is set up to *read it*.

Why AI informate is not just business intelligence in a new package

A fair objection at this point is that informate sounds like what enterprises have been calling business intelligence for the last two decades. If informate is just BI with better-looking dashboards, the transform thesis does not need a separate middle layer. The distinction is worth naming precisely, because it explains why AI opens categories of informate work that were not previously possible.

Traditional BI and analytics read *structured* data — the numbers in the fields, the transactions in the tables, the dimensions you thought to capture when you designed the data model. The information accessible through BI was the information you already knew enough to ask about. Everything else — the free text of a customer service call, the narrative of a medical record, the prose of a contract, the unscripted conversation in a Slack channel, the body of an email — was inaccessible at scale. It was data, but it was not queryable data.

AI informate changes this in three specific ways. First, large language models can read unstructured data with comprehension, not just keyword matching. Second, they can synthesize across sources rather than just retrieve from within them. Third, they can produce inference in natural language rather than rows in a table — surfacing patterns an analyst would have needed weeks of work to describe. And they can do all of this at a volume and a pace that was operationally impossible even two years ago.

The shift matters because in most organizations, the valuable information has always lived in the unstructured half of the data estate — the 80% that traditional BI could not reach. AI informate opens that half for the first time. That is not a cosmetic improvement on BI. It is a different category of capability.

Moody's: informate as strategy in B2B financial analytics

The clearest recent example of an informate-layer strategy executed well is Moody's. The credit-ratings and risk-analytics firm spent 2024-2025 building what it calls "Moody's Agentic Solutions" (MAS) — a suite of AI-powered systems in which coordinated agents perform the

component tasks of credit analysis: pulling financial statements, summarizing earnings calls, cross-referencing sector and macro data, assembling the material for a credit memo. The work that used to take an analyst several days now happens in minutes.

The superficial reading is that this is an automate story: Moody's is using AI to make credit analysts faster. That reading misses what actually happened. Moody's did not just speed up the existing workflow. It restructured the relationship between the analyst and the underlying data. The data foundation — more than 590 million entities, decades of historical ratings, structured sector and macro intelligence — had always existed, but it was accessible only in fragments, through bespoke queries that each analyst ran individually. The agentic layer made it continuously, coherently queryable by the whole organization. The analysts' job shifted from *assembling* the memo to *reviewing and interrogating* a memo the system had assembled.

That is an informate move, not an automate one. The output is not a faster version of the old credit analysis. It is a categorically different kind of credit analysis — one in which the human is freed to spend time on the questions that require judgment, and the system handles the questions that require thoroughness. Moody's has since opened the same agentic solutions to its enterprise customers through a Model Context Protocol integration with Claude, letting banks and asset managers run credit and compliance workflows natively inside their own AI environments. The business model has shifted too: Moody's is selling decision-grade intelligence as a service, not just ratings.

The Moody's case also illustrates why informate is hard to sell to a board. The automate metrics — hours saved per analyst, time to memo completion — are easy to report. The informate metrics — the quality of the questions analysts are now asking, the risks that are being surfaced earlier, the decisions that are being made with better underlying intelligence — are harder to put in a slide. But they are what actually determine whether the next generation of credit products can exist.

Mass General Brigham: the unstructured half, finally readable

A second informate case, drawn from a completely different industry, sharpens the point about unstructured data. Mass General Brigham is one of the largest academic health systems in the United States. Its clinicians generate an enormous volume of free-text clinical notes — descriptions of patient encounters, symptoms, test results, treatment responses — most of which has historically been inaccessible to anything except a human reviewer. The structured data in the health system's EHR — diagnosis codes, lab values, medication lists — is a fraction of what the clinicians actually know about their patients.

In 2023, researchers at the system's Accelerator for Clinical Transformation built RECTIFIER — a Retrieval-Augmented Generation-Enabled Clinical Trial Infrastructure for Inclusion Exclusion Review. The system reads patient records to determine eligibility for clinical trials. The traditional process was manual, labor-intensive, and inaccurate. Research coordinators spent hours per case reviewing charts, and they frequently missed eligible patients because diagnosis codes reflected billing rather than clinical reality. RECTIFIER reads the unstructured narrative — the notes, the reports, the actual clinical language — and surfaces eligible patients with a speed and accuracy manual review could not match. In one published trial, it cut melanoma trial identification from more than seven hours of manual review to 2.5 minutes.

By December 2025, Mass General Brigham had spun out the technology as a new company, AlwithCare, with plans to license the capability to other health systems. In parallel, the organization rolled out ambient clinical documentation AI to more than 3,000 clinicians and reported a 21.2% absolute reduction in physician burnout at 84 days — a measurable improvement in a variable every health system in the country is losing ground on.

The Mass General case is pure informate. The clinical narrative existed before the AI; the clinical workflow existed before the AI; the regulatory requirements existed before the AI. What changed is that a category of information previously inaccessible at scale is now queryable, inferrable, and actionable. That is a categorically different thing from automating an existing workflow or transforming the organization. It is the organization suddenly being able to see what it could not see before — and the new visibility is the foundation on which the next wave of transform-level moves will be built.

Informatize as the bridge to transform

The reason informatize matters so much for the transform layer is that organizations cannot transform what they cannot see. Every transform-level initiative begins with a moment of clarity about a part of the business that was previously opaque. A retailer realizes that its replenishment decisions are systematically biased because it can finally see, for the first time, the full causal chain between promotion, inventory, and margin. A bank realizes that its underwriting model is leaving a segment of profitable customers on the table because the data that would have surfaced them has never been assembled in one place. A manufacturer realizes that its service business could be repriced as a subscription because the telemetry from its installed base suddenly makes usage-based billing computable.

None of those moments of clarity happen without the informatize layer doing its work. The organization has to build the data infrastructure, the analytical muscle, and — this is the part that gets under-discussed — the **managerial habit of actually consuming what AI makes visible**. The AI will produce the visibility. Whether anyone in the executive team spends the time to look at it, interrogate it, and act on it is a cultural question, not a technical one.

BCG's research on AI maturity identifies this as the single most reliable predictor of whether an informatize-layer investment produces transform-layer results. The organizations that capture the value are the ones where senior executives — not analysts, not data teams, the executives themselves — spend time with the visibility the AI has produced. The organizations that do not capture the value have terrific dashboards that no one in the C-suite has ever actually sat with.

Informatize-Layer Investments Worth Naming • Analytical infrastructure that makes the data AI produces continuously queryable by the organization, not trapped in individual tools. • Managerial rituals — weekly or monthly reviews at the executive level that sit with the visibility AI has generated and decide what to do about it. • Cross-functional learning loops that let the pattern one team surfaces reach the teams that can act on it. • Data governance that makes the informatize layer trustworthy enough to act on, rather than a continuously contested set of numbers nobody believes.

VI. The Transform Layer: Where Durable Advantage Lives

Transform is the layer where the enterprise gets reshaped. A transform-layer AI initiative does not make the existing work faster; it changes what the organization is in business to do. It produces

new products, new workflows, new business models, or — in the most consequential cases — a combination of all three. It is the layer where the 20% EBITDA uplifts McKinsey documented in its April 2026 *AI Transformation Manifesto*, a study of twenty leading AI transformations, actually come from. And it is the layer where almost no enterprises in the broader population are currently working.

McKinsey's finding from that twenty-company study is worth stating precisely. The companies that executed genuine transform-layer AI initiatives produced, on average, a 20% EBITDA uplift, reached breakeven in one to two years, and generated \$3 of incremental EBITDA for every \$1 invested. They did not achieve these numbers by deploying AI broadly across the enterprise. They **concentrated** on one to three business domains, reinvented those domains, and accepted that the reinvention was the work. The disciplined narrowness was part of the strategy. Most enterprises attempting transform-layer work fail the opposite way — spreading a thin layer of AI across fifty use cases and finding no competitive advantage anywhere.

What transform actually looks like: four cases across four industries

The clearest way to explain transform is to look at what the organizations doing it are actually doing — and at what one prominent organization tried, executed at scale, and visibly failed at in a way the framework would have predicted. The cautionary case below — GE's Predix platform — establishes the failure mode against which the five success cases that follow can be read. The five winners come from industries as different as agricultural equipment, mass-market home improvement retail, biopharma, fintech infrastructure, and consumer language education. They have almost nothing in common operationally. They have everything in common strategically — and what they share is precisely what GE Predix lacked.

The cautionary case: GE Predix — what transform-layer failure looks like

Before turning to the five winners, the framework demands a case where transform was attempted, executed at scale, and visibly failed. The canonical example is GE Predix, and the failure pattern it displays is one the framework would have predicted.

Beginning around 2013, GE Chairman and CEO Jeff Immelt committed the company to a transformation from industrial conglomerate to "digital industrial" enterprise. The thesis was that GE's installed base — gas turbines, jet engines, locomotives, MRI machines, wind farms — generated vast quantities of operational data, and that the company that built the cloud platform to capture, analyze, and monetize that data would capture the value those assets had been creating without payment for a century. Predix, formally launched in 2015, was that platform. The associated business unit, GE Digital, was headquartered in San Ramon, California — deliberately distant from GE's industrial heart in Schenectady and Boston — and was projected to grow toward roughly thirty thousand engineers and approximately fifteen billion dollars in software revenue by 2020.

The narrative was coherent. The capital was committed. The CEO was personally aligned. By any measure of executive ambition this was a transform-layer move, framed as such, funded as such, and announced as such to the markets and the workforce. By late 2017 it was visibly failing. Predix revenue was a fraction of the projection, customer adoption was concentrated in a small number of friendly accounts, and the underlying platform had real performance and integration issues. Immelt was replaced by John Flannery in August 2017, who was replaced by Larry Culp in October 2018.

By 2020 the digital-industrial narrative had been quietly retired. GE Digital was reorganized and most of its initial ambitions abandoned; the company eventually broke itself into three pieces.

Read through the three-order framework, the failure was overdetermined. First, the transform leap was attempted without the informate substrate. Most of GE's industrial customers — utilities, airlines, hospitals — could not cleanly access their own machine data in the form Predix required, because the data infrastructure inside the customer organization had not been built. GE was trying to sell decision-grade analytics to organizations that were still struggling to instrument their assets and clean their data. The platform was solving a problem one layer above where the customer actually was. Second, the executive team owned the ambition but not the fluency. Immelt championed digital industrial publicly and at length, but the operating leadership remained classic industrial GE — running the businesses on the metrics and rhythms that had made the company successful when it was selling turbines, not platforms. The fluency required to evaluate Predix proposals on their actual merits was not distributed across the executive team. Third, the governance structure could not protect the long-cycle bet. GE's quarterly earnings were under continuous pressure throughout the Predix period, including from activist investor Nelson Peltz's Trian Partners, which took a board seat in 2017; the company could not sustain the intangible-investment middle of the J-curve when the financial markets demanded near-term cost cuts. Fourth, the complementary investments — reskilling the industrial salesforce into a software sales motion, restructuring customer relationships from transactional to subscription, rebuilding the operating cadence around a different set of metrics — were under-resourced relative to the platform investment itself. Capital flowed to the algorithms and infrastructure; the seventy percent of effort BCG would later describe as belonging to people and process did not.

The failure of Predix is sometimes told as a story about a platform that did not work. That telling is incomplete. Predix the technology had real flaws, but the underlying failure was not technical. It was the failure to do the work the framework names — to build the informate substrate, develop the executive fluency, protect the long-cycle investment from quarterly pressure, and fund the intangible complements at the level the transform required. The five cases that follow each succeeded by doing one or more of those things deliberately. Reading them in the absence of GE Predix would obscure the lesson. Reading GE Predix first makes the lesson visible.

Case 1: John Deere — from equipment manufacturer to data platform

John Deere has been quietly executing one of the most ambitious transform-layer AI strategies in the economy. The core product is a tractor, or a sprayer, or a combine — a large piece of machinery a farmer buys, uses for a decade or two, and eventually trades in. The traditional business is cyclical, capital-intensive, and dependent on commodity prices. Its financial markets value it as an industrial manufacturer.

Beginning with its 2017 acquisition of Blue River Technology, Deere has been reconceiving the tractor as a data platform. The most visible product is See & Spray, a computer-vision system that identifies individual weeds and activates only the spray nozzles directly above them. In 2025, See & Spray covered 5 million acres and saved farmers an estimated 31 million gallons of herbicide mix, with typical farms reducing herbicide use by 50–70%. That is a meaningful sustainability and cost story on its own. It is not the transform story.

The transform story is that the same sensors that enable See & Spray also generate a continuously accruing record of soil conditions, crop health, equipment performance, operator behavior, and field-level yield — streamed into the John Deere Operations Center. The tractor is no longer just a piece of equipment; it is a rolling data collection system. That changes what Deere sells. Under its "Leap Ambitions" strategy, the company has committed to generating 10% of revenue from recurring sources — software licenses, per-acre usage fees, subscription access to autonomy features — by 2030. That is a business model transformation, not an efficiency play. Financial markets value recurring software revenue at a dramatically higher multiple than cyclical industrial revenue, which is exactly the re-rating Deere is signaling it expects.

What makes this a transform case rather than an automate case is that Deere did not just add AI to its existing products. It used AI to change what the product *is*. A tractor is now a sensor. A sprayer is now a computer-vision system. The farm is now a platform. None of that would have been visible to a leadership team that thought about AI primarily as a way to make the factory floor more efficient.

Case 2: Home Depot — the orange apron, scaled

The second transform case is a different kind of legacy business. Home Depot operates more than 2,350 stores, employs around 470,000 people, and sits at the center of a roughly \$1.1 trillion fragmented home improvement market. Its advantage has been built, for fifty years, on the expertise of its orange-aproned associates — the people on the floor who can answer the DIY customer's real question, which is almost never about a product and almost always about a project. That expertise is the brand. It is also the constraint: capped at the number of aprons on the floor at any hour, inconsistent across stores, and unavailable when the customer is standing in their own kitchen at 9 PM with a leaking faucet and a decision to make.

Magic Apron, launched internally in October 2024 and unveiled publicly in March 2025, was framed by Home Depot explicitly around this constraint. Jordan Broggi, EVP of customer experience and president of online, described the goal at launch in a single sentence: "bring that same expertise to the digital world." The initial release was deliberately modest — conversational AI embedded on millions of product pages across homedepot.com and the mobile app, trained on Home Depot's proprietary project expertise and product catalog, answering how-to questions and summarizing customer reviews. It was a first-of-its-kind deployment in home improvement retail, but it was not yet a reconception of what the business does. It was, at that point, a better version of product discovery.

What the first fourteen months revealed is worth flagging specifically for executive audiences considering comparable moves. At its December 2025 investor day, Home Depot reported that Magic Apron had produced "consistent growth in customer adoption" and was "translating into higher levels of conversion." AI-driven search and Magic Apron together, the company told analysts, had "significantly boosted online conversion." The numbers were not disclosed in detail, but the strategic posture was revealing. Home Depot told its investors explicitly that they had gone "front-footed" on the experience change before the full returns had materialized — that the first fourteen months had been about adoption and learning, not about the payoff curve. That is the intangible middle of the J-curve, communicated deliberately to Wall Street.

The second act arrived in January 2026 at NRF, in the form of a significantly expanded Google Cloud partnership and a reconception of what Magic Apron is for. The assistant moved past product-page Q\&A. It is now conversational across the full range of home improvement projects — from fixing a leaky faucet to planning a full kitchen remodel. Multimodal capabilities (image upload, visualization) are on the roadmap. A new in-store version of Magic Apron, piloting in select locations, integrates the agentic experience with real-time store-level inventory and product locations, directing a customer to a specific aisle and bay and offering technical guidance at the point of decision. And — in the piece of the strategy most worth paying attention to — Home Depot announced it would participate in Google's agentic shopping experiences across AI Mode in Google Search and the Gemini app, extending Home Depot's reach beyond its own digital properties into whatever agentic commerce channels emerge.

The Pro-facing piece of the January expansion is the most instructive of all. Material List Builder AI, launched January 26, 2026, lets a professional contractor describe a project in natural language ("primary bathroom remodel, replacing the double vanity, adding recessed lighting, black 12-by-18 tile for the floor..."), or use voice, or paste an existing list, and receive a grouped, phase-organized material list complete with preferred pricing and current inventory availability — ready to order through a single supplier. What used to take hours of SKU searching and spreadsheet rebuilding now takes minutes. That sounds like automate — faster version of existing work — but what it actually accomplishes is strategically different. Home Depot is moving upstream in the Pro's workflow. The company is no longer just the place the Pro orders materials; it is the place the Pro plans the job. Once the plan lives inside Home Depot's Project Planning platform, the Pro's switching cost to a competitor rises substantially. The company is not automating the order; it is claiming the workflow that produces the order.

Read through the three-order framework, Home Depot is doing several things simultaneously. Parts of Magic Apron are clearly automate — faster answers, faster material lists. Parts are clearly informate — the continuously accruing record of what DIY customers are actually trying to do, where they get stuck, which products they consider, visible to Home Depot at a resolution the in-store channel never produced. And parts are transform — the migration from "retailer that sells products" to "partner that helps the customer complete the project," with an ambition, visible in the Google AI Mode and Gemini integrations, to have Home Depot's expertise present wherever the customer's planning actually happens, not only on Home Depot's own properties.

Whether the transform piece matures depends on what the next eighteen to twenty-four months produce — and, honestly, on whether Home Depot can translate project ownership into a sustainable change in its economics rather than just in its user experience. Executive observers should watch three things: conversion rates on Magic Apron-assisted sessions versus unassisted sessions, Pro Xtra adoption and repeat-order rates on Material List Builder as a signal of whether the workflow claim is holding, and the company's commentary on the relationship between its digital tools and its physical-store performance as a signal of whether it is executing interconnected retail or running two separate businesses in parallel.

A disclosed-results counterpart arrived in May 2026, when Walmart's first-quarter earnings commentary gave its agent strategy the kind of numbers Home Depot has so far withheld. Walmart told investors that weekly active users of Sparky, its AI shopping agent and the consumer-facing piece of a broader super-agent strategy, more than doubled in the quarter, that units purchased through Sparky grew more than fourfold sequentially, and that Sparky users carry an average order

value roughly 35% higher than customers who do not use the agent. The figures are company-reported rather than independently measured, and the order-value gap almost certainly carries some selection effect, since the most engaged shoppers try agents first. But the disclosure itself is the strategic signal. Walmart is now telling Wall Street that agentic commerce is producing measurable lift in order value and unit volume, which raises the bar for what analysts will eventually expect Home Depot to show for Magic Apron's still-undisclosed conversion gains.

The broader lesson for legacy-industry executives is the one that makes this case worth its pages in a whitepaper written primarily for them. Home Depot did not have to rebuild itself to run Magic Apron. The proprietary product catalog, the decade of AI at scale on the website, the in-store expertise captured in associate training materials, the Pro Xtra loyalty infrastructure, the physical store footprint — all of it was already there. The transform move was to recombine what the company already had, in the presence of a new technology, into a different customer relationship. That is what transform tends to look like in a legacy industry. It is not a replacement of the old business with a new one. It is the reorganization of existing assets around what the technology now makes possible — with the understanding that the enterprise's existing scale, trust, and data become advantages rather than drags, provided the executive team has the fluency to see them that way.

Case 3: Moderna — the shape of a biotech firm, rethought

Moderna's relationship with OpenAI, which began in 2023, started as most enterprise AI partnerships do — with a broad deployment of ChatGPT Enterprise and an internal chatbot called mChat that had reached 80% adoption within months. What has happened since is something different. Moderna has built more than 3,000 internal GPTs — specialized AI agents for particular scientific, operational, and administrative tasks. Its legal team reports 100% ChatGPT Enterprise adoption. Its Dose ID GPT applies standard selection criteria to clinical trial dose evaluation, providing rationale, citing sources, and giving the human reviewer a structured starting point that did not previously exist.

The most telling structural move came in 2024, when Moderna merged its HR and digital technology functions under a single "chief people and digital technology officer." That is not a governance tweak. It is a statement that the workforce and the AI toolkit are now the same problem — that you cannot redesign one without redesigning the other.

The language Moderna CEO Stéphane Bancel has used publicly makes the transform thesis explicit: "If we had to do it the old biopharmaceutical ways, we might need a hundred thousand people today. We really believe we can maximize our impact on patients with a few thousand people, using technology and AI to scale the company." That is a specific, falsifiable claim about what Moderna is becoming — a differently-shaped biopharma company, not a conventional biopharma with more AI inside it. The claim may or may not prove correct commercially. But it is the kind of narrative that can organize an executive team, a board, and a workforce around a coherent direction, and it is the kind of narrative that most enterprises have not yet found for themselves.

Case 4: Ramp — institutional AI as an organizational capability

Ramp, the corporate finance software company, offers a smaller but sharper example of what transform looks like operationally. Ramp reports that 99% of its employees use AI daily. Eighty-four

percent use AI coding agents weekly. Non-engineers regularly ship production code. The numbers are striking, but the interesting question is how Ramp got there.

CEO Eric Glyman's diagnosis was that early adoption had high ceilings but shallow floors — a minority of employees were getting extraordinary value from AI while the majority were stuck on setup, configuration, and the absence of institutional context. Ramp's response was to build an internal AI system called Glass: a single, role-aware entry point to AI that comes pre-configured on install, authenticates through Okta SSO, and gives every employee immediate access to every tool the company uses. Layered on top of Glass is a skill marketplace called Dojo, where employees package their best workflows into reusable files the whole company can use; over 350 such skills have been shared internally. An AI guide called Sensei looks at an employee's role, connected tools, and recent activity, and surfaces the handful of skills most likely to be useful.

None of this is individually novel. What is novel is that Ramp built the infrastructure of institutional AI use — the shared muscle that turns individual AI productivity into organizational AI productivity. It is the explicit counter to what venture investor George Sivulka has called the "productive individuals do not make productive firms" problem: the observation that AI has made knowledge workers roughly ten times more productive without making most of their companies any more valuable, because the individual productivity has not translated into organizational productivity. Ramp's Glass is one of the first clean examples of an organization doing the translation deliberately.

Case 5: Duolingo — transform at the edge of the organization's absorption capacity

Duolingo offers the final case, and a cautionary one. In April 2025, CEO Luis von Ahn announced an "AI-first" strategy and committed to the gradual replacement of human contractors with AI for content creation. The company followed within days by launching 148 new language courses built primarily by generative AI — doubling its course catalog. For context: Duolingo's first 100 courses took twelve years to build. The next 148 took about one.

On the content-creation dimension, this is transform. The addressable product universe expanded by an order of magnitude in under a year. Duolingo can now launch a course for a language combination that would never have cleared the cost threshold under the old model. That opens a business model possibility — deeper tier-based monetization, premium AI-enabled features in the Max subscription — that did not exist at the old scale.

What Duolingo illustrates, though, is what happens when a transform narrative outruns the organization's ability to absorb it. The announcement produced a visible user backlash — complaints of course quality regression, cultural criticism of "AI-first, customers-last," social media videos of users deleting the app. The company did not reverse the strategy, but it tempered the communication significantly. The lesson is not that Duolingo was wrong to pursue transform. The lesson is that transform-layer moves require congruence across strategy, culture, communication, and customer experience — and an organization that moves one of those faster than the others pays a visible price. Duolingo's transform is real. Its execution of the surrounding change management was not yet ready for the speed of the product change.

What the winners share — and what GE Predix did not

John Deere, Home Depot, Moderna, Ramp, and Duolingo look nothing alike as businesses — industrial agriculture, mass-market home improvement retail, biotech, fintech infrastructure, and consumer education. Three are decades-old legacy enterprises with workforces in the hundreds of thousands; two are digital natives. The common thread is not what they sell. It is how they think. In each case, the leadership team did not ask "where can we apply AI to what we already do?" They asked "what could we be that we could not be before AI?" — and then built backward from the answer.

What the five winners share is precisely what GE Predix lacked. Each of them built informate substrate before scaling the transform initiative; each had an executive team whose fluency tracked the ambition; each had a governance structure that protected the long-cycle bet from short-cycle pressure; each funded the complementary investments — workforce, customer relationship, operating cadence — in proportion to the platform investment. None of these conditions is accidental. They are the four conditions the framework predicts must be present for a transform initiative to mature, and the absence of any single one of them produces the Predix outcome regardless of how strong the other three are. The five winners are five different demonstrations of the same disciplined execution. GE Predix is the demonstration of what happens when ambition runs ahead of the substrate that has to support it.

That is the transform question. It is not a question a CIO can answer. It is not a question a consulting engagement can answer. It is a question the executive team has to personally own, using personal fluency in AI capability as part of the answering process, because the answer depends on judgments about what AI can actually do that no one outside the executive team can responsibly make.

Which is why the next section is about the C-suite.

Part III — The C-Suite Implication

VII. Why the Transform Layer Cannot Be Delegated

The automate layer can be delegated. A CEO can reasonably hand an automate-layer AI portfolio to a CIO, a COO, or a chief of staff, approve the capital, set the review cadence, and stay at a comfortable distance from the day-to-day. The informate layer is harder to delegate but still possible — a chief data officer or a chief digital officer can run most of the organizational machinery, with the executive team engaging mainly at the governance and decision-making level.

The transform layer cannot be delegated. Not because it is technical — most of the hard decisions at the transform layer are not technical — but because it requires judgments about what the enterprise should become, and those judgments belong, institutionally and legally, to the chief executive and the senior team. It also requires the institutional standing to defend the J-curve dip described in Section III to boards and analysts who will pressure the firm to abandon the transform investment before it matures. No functional leader has that standing. The defense — strategic, communicative, financial — is executive work.

The delegation trap

The dynamic plays out consistently across the organizations I have worked with over the last two years. A CEO inherits an AI agenda — usually a list of pilots assembled by functional leaders, each optimizing for something measurable within their own domain. The list is long. The board wants to see progress. The executive team wants budget. The CEO, understandably, delegates the orchestration downward. A steering committee is formed. A chief AI officer is appointed, or the mandate is given to an existing CIO or CDO. The CEO stays close at a governance level but steps back from the substance.

Eighteen months later, the CEO has a portfolio of pilots, a dashboard of adoption metrics, and no clear answer to the question her board has started asking: *what are we becoming?* The steering committee cannot answer that question because the answer cuts across every function, and no single functional leader has the authority to redefine what the enterprise does. The chief AI officer cannot answer it because the question is not really about AI; it is about strategy. Only the executive team, working together, can answer it. And answering it requires a level of direct engagement with the technology most of the executive team has not yet done.

This is what BCG's 2026 AI Radar captured in its segmentation of CEOs. Seventy-two percent of CEOs now describe themselves as the primary decision-maker on AI in their organization — double the share from a year earlier. But the same survey splits those CEOs into three archetypes: Followers (about 15%), Pragmatists (about 70%), and Trailblazers (about 15%). The Trailblazers are doing measurably different work. They spend more than eight hours a week on their own AI upskilling. They allocate sixty percent of their AI budgets to workforce reskilling. They can describe, in their own words, what their company is becoming.

The gap between owning an agenda and owning a narrative is the quiet crisis of AI leadership today. Most CEOs have agreed, in the form of budget allocations and press interviews, to own their

company's AI strategy. What most of them have not yet done is the kind of personal, direct engagement with the technology that would let them tell, from their own judgment, which transform-layer bets are worth making and which are not.

What fluency actually means

AI fluency for a C-suite executive is not a technical concept. It does not require the ability to fine-tune a model, evaluate a model's architecture, or write a prompt using any particular framework. What it requires is enough direct, personal experience with the current generation of frontier AI tools that the executive can:

- **Tell, from direct use, what AI can and cannot reliably do today.** The capabilities change monthly. Executives who are not using the tools themselves are making strategic decisions based on outdated assumptions — usually assumptions that are six to twelve months behind the state of the art.
- **Smell the difference between a real transform proposal and a dressed-up automate proposal.** The dressed-up version is everywhere. It is the automate project with a transform-sized budget and a transform-sized slide deck. A fluent executive can tell the difference in under five minutes. A non-fluent executive cannot.
- **Push back on vendors and consultants with specific rather than generic questions.** The specific questions surface real capability gaps. The generic questions produce generic answers.
- **Model the use of AI for the organization.** The fastest way to signal that AI is a strategic priority is for the executive team to use it visibly and routinely — in board preparation, in strategic analysis, in communication. Employees notice. Middle managers notice. The culture shifts faster on the visible practice of the executive team than on any memo they issue.

Role-specific fluency requirements

The fluency each executive needs is not identical. It varies with the role. The table below sketches what transform-era fluency looks like for each of the major C-suite seats.

Role	What transform-era fluency looks like
CEO	Personal AI practice of 5–10 hours/week. Ability to articulate, in own words, what the company is becoming. Board-level conversation about AI as strategic posture rather than tool selection.
CFO	Understanding of the J-curve economics of AI investment. Capital-allocation framework that protects transform budgets from automate-urgency pressure. Fluency in the difference between short-cycle ROI (automate) and long-cycle value creation (transform).
CHRO	Direct experience using AI in HR workflows. Workforce redesign capability — not just reskilling programs, but the actual redesign of roles and teams around AI. Change management muscle for the sustained organizational transition transform requires.
COO / CIO	Personal fluency in the tooling. Governance design that lets the organization move fast on automate without compromising informate data quality. Infrastructure strategy that supports all three layers, not just the bottom one.
CMO	Direct experience using AI in customer-facing workflows. Judgment about where AI augments the customer relationship and where it degrades it. Ability to distinguish personalization that improves experience from personalization that feels intrusive.
General Counsel	Personal engagement with AI governance questions. Judgment about risk posture — where the organization should be cautious, where it should be bold, and what legal guardrails support transform without stifling it.

Role	What transform-era fluency looks like
Board	Enough collective fluency to ask substantive questions rather than generic ones. An AI literacy program for the board itself, not just the management team.

The common thread across every row is direct, personal engagement with the technology. A CFO who has never personally used a frontier AI tool to do real work cannot meaningfully evaluate a transform-layer capital request. A CHRO who has not used AI in HR workflows cannot credibly lead the workforce redesign the transform will require. A general counsel who has not engaged with the governance questions directly cannot write the guardrails the organization needs. The delegation downward — "our CIO handles AI for us" — is the organizational equivalent of a CEO delegating strategy.

The cost of delegation

The cost of non-fluency at the C-suite level is not visible immediately. It shows up downstream, in three predictable forms.

First, the organization spends heavily on automate-layer initiatives that produce real near-term value and no durable advantage, because the C-suite cannot evaluate the transform-layer alternatives on their merits.

Second, the organization's AI governance becomes reactive rather than strategic — focused on what to prevent rather than what to enable — because the executive team lacks the personal judgment to set a posture that is both safe and ambitious.

Third, the organization loses its most AI-capable talent. High-capability AI employees do not stay in organizations where the executive team does not understand what they do. They leave for organizations where the strategic conversation is happening at a level that matches their competence. By the time the executive team realizes the talent attrition has happened, the window to execute transform has usually closed.

These three costs compound. An organization that loses its AI talent this year will struggle to fund a transform initiative next year, because the internal capability to execute will not be there. An organization that spends all its AI budget on automate will have less runway for transform when the time comes, because the intangible middle of the J-curve has to be funded, and automate initiatives do not generate the complementary assets the transform layer requires. Delegation is not a neutral choice. It is a choice that forecloses options over time.

VIII. The CEO Playbook: What to Do in the Next Ninety Days

The previous section described what the C-suite should no longer delegate. This section describes what the C-suite should actually *do* — specifically, in the next ninety days.

The playbook has four parts: a portfolio audit, a personal practice regimen, a governance structure that protects the transform agenda, and a first transform domain chosen and funded with discipline. None of the four is technically difficult. All four require the executive team to actually do the work, together, rather than handing it off.

Before the playbook: a diagnostic for classifying any AI proposal

Before running the playbook, the executive team needs a shared vocabulary for classifying proposals. The five questions below, applied to any AI initiative on the current portfolio or any new one coming across the desk, produce a quick and usable classification into automate, informate, or transform. The diagnostic is not a substitute for detailed business-case review; it is a filter that exposes the strategic category of an investment before the financial case obscures it.

The Five-Question Diagnostic 1. Does this change what we do, or how fast we do what we already do? "*Faster*" → *automate*. "*Different*" → *informate or transform*. 2. Is the ROI calculable today from hours saved, headcount avoided, or costs reduced? *Yes* → *automate*. *No* — *the value depends on new revenue, new capability, or new customer behavior* → *informate or transform*. 3. Could a competitor replicate this within 12 months using off-the-shelf AI tools? *Yes* → *automate*. *No* — *the capability depends on proprietary data, workflow redesign, or organizational change competitors cannot shortcut* → *informate or transform*. 4. Does executing this require cross-functional authority to redesign work? *No* — *one function can execute within its existing mandate* → *automate*. *Yes* — *it crosses functional lines by definition* → *transform*. 5. Will the customer experience or product itself be fundamentally different after, or just faster? *Same experience, faster* → *automate*. *Better-informed decisions about the same experience* → *informate*. *Different experience or different product* → *transform*. Scoring: Mostly "*faster / bounded / replicable / single-function / same-faster*" answers indicate an automate project. Mostly "*different / long-cycle / hard-to-replicate / cross-functional / different*" answers indicate transform. Mixed answers with strong visibility-and-decision-quality themes indicate informate. Watch for: Proposals that answer "*different*" on questions 1 and 5 but "*calculable today*" on question 2, and "*within one function*" on question 4. This is the signature pattern of automate wearing transform clothing — a common category that deserves closer scrutiny before it is funded at a transform budget level.

1. Audit the current AI portfolio across the three layers

Begin with a simple exercise. Take the current portfolio of AI initiatives — every pilot, every deployment, every committed budget — and assign each to one of the three layers. Do the assignment as an executive team, in a room, not by memo. The disagreements will be instructive.

Most executive teams doing this exercise for the first time discover the same pattern: 70–85% of the portfolio is automate, 10–20% is informate, and the transform allocation is effectively zero. Initiatives that were presented as transform turn out, on examination, to be automate projects wearing transform language. This is not a scandal. It is the natural consequence of building an AI portfolio bottom-up through functional leaders, each of whom is optimizing for measurable outcomes in their own domain.

The goal of the audit is not to eliminate automate projects. It is to expose the portfolio composition and force an explicit conversation about whether the composition matches the strategic ambition. If the organization's ambition is to be a durable transform-layer winner in its industry, a 0% transform allocation is a policy to choose, not a condition to accept.

A simple portfolio target for 2026–2027 Automate: 50–60% of AI investment. Funds near-term value and credibility. Savings reinvested, not swept. Informate: 25–30% of AI investment. Funds the data infrastructure, analytical muscle, and managerial habits that make transform possible. Transform: 15–25% of AI investment, concentrated in one to three business domains. Funds the long-cycle bets that produce durable advantage. *The specific percentages matter less than the principle: transform must have a real allocation, protected from automate urgency, and concentrated rather than sprayed across the enterprise.*

2. Establish a personal practice regimen

Every member of the executive team should personally use frontier AI tools for real work — not demo work, not training exercises, real work — for at least five hours a week. The BCG Trailblazer

finding of eight hours a week is a better target for the CEO. The point is not hours. The point is that direct, repeated use is the only way to develop the judgment the role now requires.

Concretely, this means each executive should identify three to five of their own recurring work tasks — tasks they already do every week — and move some portion of them into an AI-assisted workflow. Board preparation. Industry analysis. Strategic memos. Meeting summaries. Communication drafting. Competitor research. The tasks are less important than the principle: the executive should have direct, personal experience of what AI does well, what it does poorly, and how the work actually changes when AI is in the loop.

This is not reskilling. It is the minimum condition for making transform-level strategic decisions. An executive who has never used a frontier AI tool for an hour of actual work is in the same position as an executive in 2002 who had never used email. They can delegate the use to others; they cannot delegate the strategic intuition.

3. Build the governance that protects transform from automate urgency

The single most common way transform initiatives die is not that they fail on their merits. It is that they get starved. A transform project is a long-cycle investment — 18 to 36 months to real value — competing for attention against automate projects that produce measurable savings in one or two quarters. In almost every quarterly review, the automate project looks better. Over eight quarters of reviews, the transform project gets cut.

The governance structure has to recognize this failure mode explicitly. Three practical moves:

- **Separate the review cycle.** Automate projects should be reviewed on a quarterly cadence against short-cycle metrics. Transform projects should be reviewed on an annual cadence against milestone-based progress toward a clearly articulated future state. Reviewing them against the same metrics on the same cadence kills the transform projects.
- **Fence the transform budget.** Automate savings should not be usable to fund other automate savings. They should flow, by policy, into the informate and transform layers. The CFO's job is to protect this flow from the natural organizational pressure to return savings to the P&L.
- **Assign ownership to the executive team, not a committee.** Each transform initiative should have a named member of the executive team as its owner, with a direct reporting line to the CEO. Steering committees are where transform initiatives go to die by consensus.

4. Pick a first transform domain and fund it with discipline

The McKinsey finding on transform-layer success is that the winners concentrate. One to three domains, reinvented. Not fifty use cases. The first transform domain choice is a capital-allocation and competitive-position decision, and it belongs to the CEO.

Three criteria separate a good first transform domain from a bad one. First, the domain should be strategically meaningful — a change in how the domain operates should produce a visible shift in the enterprise's competitive position, not just a line-item improvement in a P&L. Second, the organization should have, or be willing to build, the informate infrastructure the transform requires. Transforming a domain you cannot see into is attempting the second floor of a house before the first floor exists. Third, the domain should have a named executive owner with the

authority to redesign work across functional lines. Transform by definition crosses functions; an owner without cross-functional authority cannot execute.

For most enterprises, the first transform domain is customer experience, product development, risk analysis, or the service layer. For industrial businesses, it is often the intersection of a physical product and the data that product generates. For services businesses, it is often the professional workflow at the core of the offering — the audit, the credit memo, the legal brief, the underwriting decision. The specific answer varies. What does not vary is that the choice has to be made, explicitly, by the executive team — and then protected from the automate-urgency pressure that will try to pull the budget back toward near-term savings.

The first-quarter checklist for the executive team

- **Conduct the portfolio audit across the three layers.** Expose current composition. Set a target for next quarter.
- **Establish personal practice.** Every executive identifies three to five recurring tasks to move into AI-assisted workflows. The CEO targets eight hours a week of direct use.
- **Fence the transform budget.** Automate savings flow to informate and transform, not back to the P\&L. Transform reviews on an annual milestone cadence.
- **Choose the first transform domain.** Pick one. Name an executive owner. Fund it for 18 months without quarterly re-justification.
- **Commit to a collective learning cadence.** Executive team spends two to three hours a month together on what the AI work is teaching the enterprise — not a status report, an actual learning conversation.

IX. The CFO Playbook: Defending the J-Curve

Section VIII organized what the CEO and the executive team must do to launch the transform agenda. This section organizes what the CFO specifically must do — and why, in the AI era, those moves are not optional add-ons to the existing CFO agenda but the decisive input to whether the rest of the playbook survives. As Section III argued, the intangible middle of the J-curve is where most transform initiatives die; the CFO is the executive who has to defend the dip in the language the boardroom understands. The four moves below describe what transform-era finance leadership actually requires.

1. Model the J-curve explicitly, and communicate it in advance

Analysts, boards, and internal stakeholders will not sit patiently through an unexplained dip in AI-related margins. The CFO's first job is to make the J-curve shape explicit in financial communications before the dip arrives. That means explaining — in capital-markets language — why a transform investment produces intangible complementary assets that are not yet capitalized, why the reported productivity numbers will lag the real productivity numbers by 12 to 24 months, and why the right comparison for transform spend is not last year's P\&L but the cost of not making the investment at all. The enterprises that get this wrong do not lose the transform bet on its merits. They lose it in the quarterly earnings call three quarters in, when the language was not prepared.

2. Establish a layer-specific KPI framework

One of the most common CFO mistakes in AI investment is applying automate-layer KPIs (cycle time, headcount avoidance, cost per unit) to informate and transform initiatives. The KPIs do not work, the initiatives miss their targets by the wrong measure, and they get killed. The CFO's framework needs to carry three distinct KPI sets, applied to the right layer.

Layer	What to measure
Automate	Cycle time reduction; cost per unit of work; deflection rate; error rate; payback period (typically < 12 months); quality floor vs. baseline.
Informate	Decision quality indicators (rework rate, decision latency, coverage of analyzed population); executive engagement with AI-generated visibility (time spent, decisions made); data quality and trust (completeness, timeliness, source reliability).
Transform	Milestone progress toward future state; customer adoption of new capabilities; new revenue attributable to transform-enabled products; workforce capability metrics; competitive position indicators (category creation, re-rating by analysts, talent inflow).

3. Fence the transform budget from automate urgency

Automate savings produce continuous pressure, every quarter, to be returned to the P&L. The pressure comes from investors, from analysts, from board members, and often from the CEO. The CFO's job — the specifically CFO job — is to resist that pressure on behalf of the transform layer. That requires three concrete moves: a written policy that automate savings above a threshold flow to informate and transform by default; a separate review cycle for transform initiatives (annual, milestone-based) that does not share a calendar with automate reviews; and a governance seat in which the CFO is the named defender of the long-cycle investments when quarterly pressure builds.

4. Know when to say no

The CFO's most consequential AI judgment is often not about approving an investment. It is about declining one. Four patterns deserve a default "no" until they are restructured: automate proposals without a quality floor (the Klarna mistake); informate proposals where no executive has committed to actually using the visibility the investment will produce; transform proposals that do not have cross-functional authority named in the business case; and any proposal whose ROI depends on the assumption that "AI will keep improving so the numbers will work out." That last category is the one most likely to be approved because it sounds sophisticated. It is also the one most likely to produce the failed initiatives that make future AI investment harder to fund.

The CFO who runs this four-part agenda becomes the enterprise's defender of long-cycle value creation in a period that will relentlessly pull the organization toward short-cycle metrics. That role is not a compliance function or a budgetary one. It is strategic leadership — and in the AI era, it is as close to decisive as any role in the C-suite other than the CEO's.

X. Answering the Hard Questions

The argument of this paper will provoke at least five objections in any honest executive reader. This section names the strongest of them and engages each directly. The transform thesis does not require that these objections be ignored. It requires that they be answered.

Objection 1: Transform is a luxury we cannot afford while competitors use automate to eat our margins

"My industry is being re-priced right now. My competitors are cutting 20–30% out of their cost structure with AI. If I spend on transform, I lose on automate — and the automate loss is immediate."

This is the objection that decides most AI strategies today, and it is also the one most likely to lead to the wrong answer. The implicit assumption is that automate and transform are competing uses of budget — that every dollar spent on a transform initiative is a dollar not spent on an automate initiative that would produce near-term savings.

The assumption is wrong in both directions. First, the automate savings themselves are the fuel for the transform budget. An automate-first strategy that reinvests the savings into informate and transform layers produces both near-term P&L improvement *and* long-term capability. An automate-first strategy that sweeps the savings to the bottom line produces one year of reported gains and three years of erosion as competitors catch up.

Second, the enterprises pulling ahead are not choosing between automate and transform. They are doing both, with a disciplined portfolio mix — roughly 50–60% automate, 25–30% informate, 15–25% transform — that lets the automate layer fund the rest. A 15% transform allocation, concentrated in one to three domains, is not a luxury. It is the minimum investment in long-run competitive position that an enterprise in any AI-exposed industry can responsibly make. The organizations that treat transform as a luxury will discover, sometime in 2027, that the organizations that treated it as table stakes have moved into positions that cannot be caught from behind.

Objection 2: The J-curve sounds like a request for permission to lose money indefinitely

"Every transform pitch I have seen rests on the argument that returns will come later. How do I know they are coming? How long is the patience window?"

A reasonable concern. Transform initiatives do go through a period where the inputs are visible and the financial outputs are not, and that period will naturally invite the question of when the curve will actually rise.

The answer is that the J-curve has milestones, not just faith. A well-run transform initiative has measurable intermediate outcomes long before it has measurable financial returns: customer adoption of new capabilities, data infrastructure built and validated, workforce capabilities assessed and strengthened, pilots expanded from concept to production deployment, and — critically — the specific organizational muscles developed that the transform requires. The CFO's job is not to defer measurement for 36 months. It is to measure the right things at the right time, and to insist that transform initiatives produce their intermediate outcomes on schedule even when the financial outcomes are still ahead.

The discipline is also in the concentration. Fifty transform pilots all failing to produce near-term financial returns is a failed AI strategy. One to three transform initiatives, each with a named executive owner and milestone-based review, producing the intermediate outcomes before the

financial ones, is how transform is supposed to look. The J-curve is not a defense of undisciplined spending. It is the framework for modeling the right kind of disciplined investment.

Objection 3: My board will not fund a 36-month bet

"I can get capital for a new plant. I can get capital for an acquisition. I cannot get capital for an AI program that will not produce measurable financial returns for three years."

Most boards have been funding 36-month bets in plant, equipment, R&D, and market expansion for their entire existence. The framing problem with AI is that it has been categorized as an IT expense rather than a capability investment, and IT expenses are evaluated on a different time horizon than R&D or infrastructure.

Reframe the conversation. The work a transform initiative does — building data infrastructure, developing organizational capabilities, reshaping workflows, reconceiving products — is capability investment. It deserves the same time horizon, the same governance posture, and the same board-level conversation that other capability investments receive. Walk the board through the J-curve explicitly. Show the distinction between automate ROI (near-term), informate outcomes (intermediate), and transform returns (long-term). Bring named cases from comparable industries into the room. Boards that have funded digital transformation will understand this framing. Boards that have not will need to be educated, and the CEO and CFO are the only people who can do that educating credibly.

Objection 4: We are a regulated industry and cannot move fast on transform

"We are a bank / insurer / healthcare system / pharmaceutical company. Our regulators will not let us do the kind of things the transform cases in this paper describe."

Regulation is a constraint. It is not the constraint. The most prominent cases in this paper — Moody's in financial analytics, Mass General Brigham in healthcare, Moderna in biopharma — are all in highly regulated industries. Their transform initiatives are advancing not because they ignored the regulators, but because they designed their AI systems with explainability, audit trails, and human oversight built in from the start.

Regulated industries actually have more to gain from transform, because the regulatory overhead that competitors face becomes a moat if you can navigate it. A bank that deploys a transform-level capability with proper compliance architecture has a structural advantage over every bank that cannot. The question is not whether regulation allows transform. It is whether the executive team has the technical literacy to design transform initiatives that satisfy the regulators — which returns the conversation to the case for C-suite fluency made in the previous section.

Objection 5: We do not have the AI talent to execute transform

"Every organization we benchmark against is hiring. We cannot compete for the talent we need to do any of this credibly."

Every organization says this. The talent problem is usually misdiagnosed.

The talent transform requires is not primarily AI engineering talent — though some is required. It is cross-functional leadership talent, with enough AI fluency to make transform-level decisions and enough organizational authority to execute them. That shortage is real, but it is a leadership

development problem, not a hiring problem. Internal upskilling, disciplined external hires as anchors in the functions where the gap is widest, and a narrow focus on one to three transform domains make the talent problem tractable at most scales.

The mistake most enterprises make is trying to solve the talent problem by hiring data scientists and ML engineers into central AI teams and then waiting for transform to emerge from those teams. It will not. Transform emerges at the intersection of business domain expertise and AI capability, which means it requires line-of-business leaders who have developed the fluency to see the opportunity themselves. That population is much smaller than the AI engineer shortage would suggest, and a much more addressable one — because the leaders are already inside the organization, waiting to be given the tools, the mandate, and the authority to lead the redesign.

XI. Conclusion: The Next Twenty-Four Months

The argument of this paper can be compressed into a single observation. The AI era will not reward enterprises that use the technology to do less of what they already do. It will reward enterprises that use it to do what they could not do before.

That asymmetry is not yet visible in the financial statements of most enterprises, because we are still in the intangible middle of the productivity J-curve. The complementary investments — process redesign, workforce reskilling, new product development, business model reinvention — are being made by a minority of firms and not being made by the majority. The majority's decision, mostly implicit, has been to capture the automate-layer gains and wait to see what happens. That decision will look defensible for another year. It will stop looking defensible sometime in 2027, when the firms that made the transform-layer investments begin to report the EBITDA numbers McKinsey's twenty-company study already documents — 20% uplifts, breakeven in one to two years, \$3 of incremental EBITDA per \$1 invested — at scale.

At that point, the argument about whether AI is transformative or merely efficient will be over. What will remain is a sorting between the firms that did the transform work in time and the firms that did not. The sorting is happening now, quietly, in the portfolio-allocation decisions and the personal-practice decisions being made in executive teams this quarter.

What the C-suite actually has to do

The work this moment requires of a chief executive is not a two-year AI roadmap. Roadmaps written today will be obsolete by the second quarter of their execution. The work is to build an organization that can absorb each quarter's changes in AI capability and convert them into sustained competitive position — an organization that moves up the automate / informate / transform stack at a cadence that matches the evolution of the technology.

That organizational work is what a chief executive has been trained for. It is leadership work. It is strategy work. It is culture work. It does not require the CEO to become a technologist. It requires the CEO to do the thing she has always done — set the direction, build the team, protect the long-cycle bets, hold the organization to a coherent narrative about what it is becoming — applied to the most consequential technology moment of her career.

The framework this paper has offered — automate to free capacity, informate to learn what the new capacity enables, transform to change what the enterprise is in business to do — is not a new idea.

It is a forty-year-old insight from a Harvard ethnographer, updated for a technology she could not have anticipated. The idea has survived this long because it names something true about how organizations absorb general-purpose technologies. The firms that treated electricity as a faster steam engine lost a generation to the firms that redesigned the factory. The firms that treat AI as faster human labor will lose a generation to the firms that redesign the work.

The difference is that this cycle will not take a generation. It will take twenty-four months.

A closing provocation

If there is one question worth taking into the next executive meeting, it is this: ***Can you describe, in your own words and without a slide, what your company is becoming as a result of AI?*** If the answer is yes, and the answer survives contact with the evidence of what the organization is actually doing, you are ahead of most of your peers. If the answer is no, or if the description is a general statement about "being AI-first" or "leveraging AI for competitive advantage," the organization is in the steam-powered electric factory — using the new technology to do the old work, with no clear picture of the redesign the technology actually makes possible.

The redesign is the work. The redesign is not optional. The only question is whether the executive team will own it personally, starting now, or discover in 2027 that the firms that did own it personally have moved into a competitive position the rest of the field cannot reach.

That is the transform imperative.

A Note on Method

This paper is itself an example of the collaborative human-AI workflow it advocates. AI handled first-pass research synthesis and drafting; the author provided strategy, framework, editorial judgment, institutional knowledge, and accountability for the final content. All framing, arguments, and conclusions reflect the author's professional judgment informed by a decade of digital transformation research and two years of direct advisory work with chief executives on AI strategy.

Copyright, reuse, and citation

© 2026 Gerald C. Kane. This paper is licensed under the Creative Commons Attribution-NonCommercial 4.0 International License (CC BY-NC 4.0), <https://creativecommons.org/licenses/by-nc/4.0/>.

You may share and adapt it for noncommercial purposes, provided you credit Gerald C. Kane as the author, link to the license, and indicate any changes you made. You may not present this material, or adaptations of it, as your own work, and you may not sell it or use it in paid training or consulting without written permission.

Cite as: Kane, G. C. (2026). *The Transform Imperative*. <https://www.profkane.com/papers/transform-imperative.pdf>
Permissions and commercial use: jerry@profkane.com