THE DIGITAL ECONOMIST

Date of Publication | June 23, 2026

Expert Insights

Sandy Carter

The Enterprise AI Culture Playbook

A Three-Pillar Framework for the 6% Who Win

CULTURE FIRST | AI READINESS | BEYOND MODELS

The Digital Economist

Email: info@thedigitaleconomist.com
Website: www.thedigitaleconomist.com

Table of Contents

Executive Summary

Boston Consulting Group has put a number on the real driver of enterprise AI success, and the number is uncomfortable for most boards (BCG 2025a).

Algorithms account for 10 percent of the work involved in AI transformation. Technology and data infrastructure account for 20 percent while the remaining 70 percent comes from people and processes. People include training and enablement, employee involvement, trust building, and change management. Processes include workflow redesign, governance, operating models, and performance measurement. They determine how work is redesigned around AI, the guardrails that keep deployments on track, and the metrics that indicate whether value is actually being created.

BCG has used this same distribution since the digital transformation era, long before generative AI existed. That 70 percent is the deciding variable in whether an enterprise generates measurable AI value, yet it is also the area in which most organizations systematically underinvest.

This paper introduces a three-pillar framework for closing the gap between AI investment and AI returns, drawing on more than a decade of work at the convergence of AI and enterprise transformation (Carter 2025). The culture framework rests on change management, data foundation, and business outcome discipline, with each pillar anchored by a global enterprise that demonstrates disciplined execution at scale.

Organizations like Qualcomm anchor the change management pillar, JPMorgan Chase anchors the data foundation pillar, and Walmart anchors the business outcome pillar. Whether measured by McKinsey's finding that only 6 percent of organizations report measurable financial returns from AI or by BCG's estimate that just 5 percent are truly built for the future, the leaders consistently occupy the same narrow band. They also share the same trait: they invested in the 70 percent.

The next twelve months will widen the gap between the 5 percent of companies generating substantial AI value and the 60 percent already classified as laggards by Boston Consulting Group (BCG 2025b). The gap widens because early advantages compound: each cycle of deployment deepens organizational learning, matures the data foundation, and refines the operating model, and every gain accelerates the next, so leaders pull away while laggards fall further behind. Companies that close the gap will do so by rebuilding the culture, data infrastructure, and operating discipline that turn AI into measurable business outcomes.

> "AI operates entirely based on the data it has to work with, and it's a system of probabilities. All AI models are trained on the public internet as well as other sources; just because something is high probability doesn't mean it's factually correct or appropriate to your specific situation. The more data of your own you provide to AI, the better it will tend to behave and the less it will hallucinate (which is tech speak for making things up)."

Christopher Penn

1.

The Enterprise AI Investment Gap

The headline numbers on enterprise AI adoption are sobering when considered together, and they tell a consistent story across every major research firm tracking the field. McKinsey reported in 2025 that 88 percent of organizations now actively use AI tools, yet only 6 percent report measurable financial results from their AI investments (McKinsey & Company 2025). The remaining 94 percent are spending real capital on initiatives they cannot tie to revenue, cost reduction, or competitive advantage, which is a structural problem of the first order.

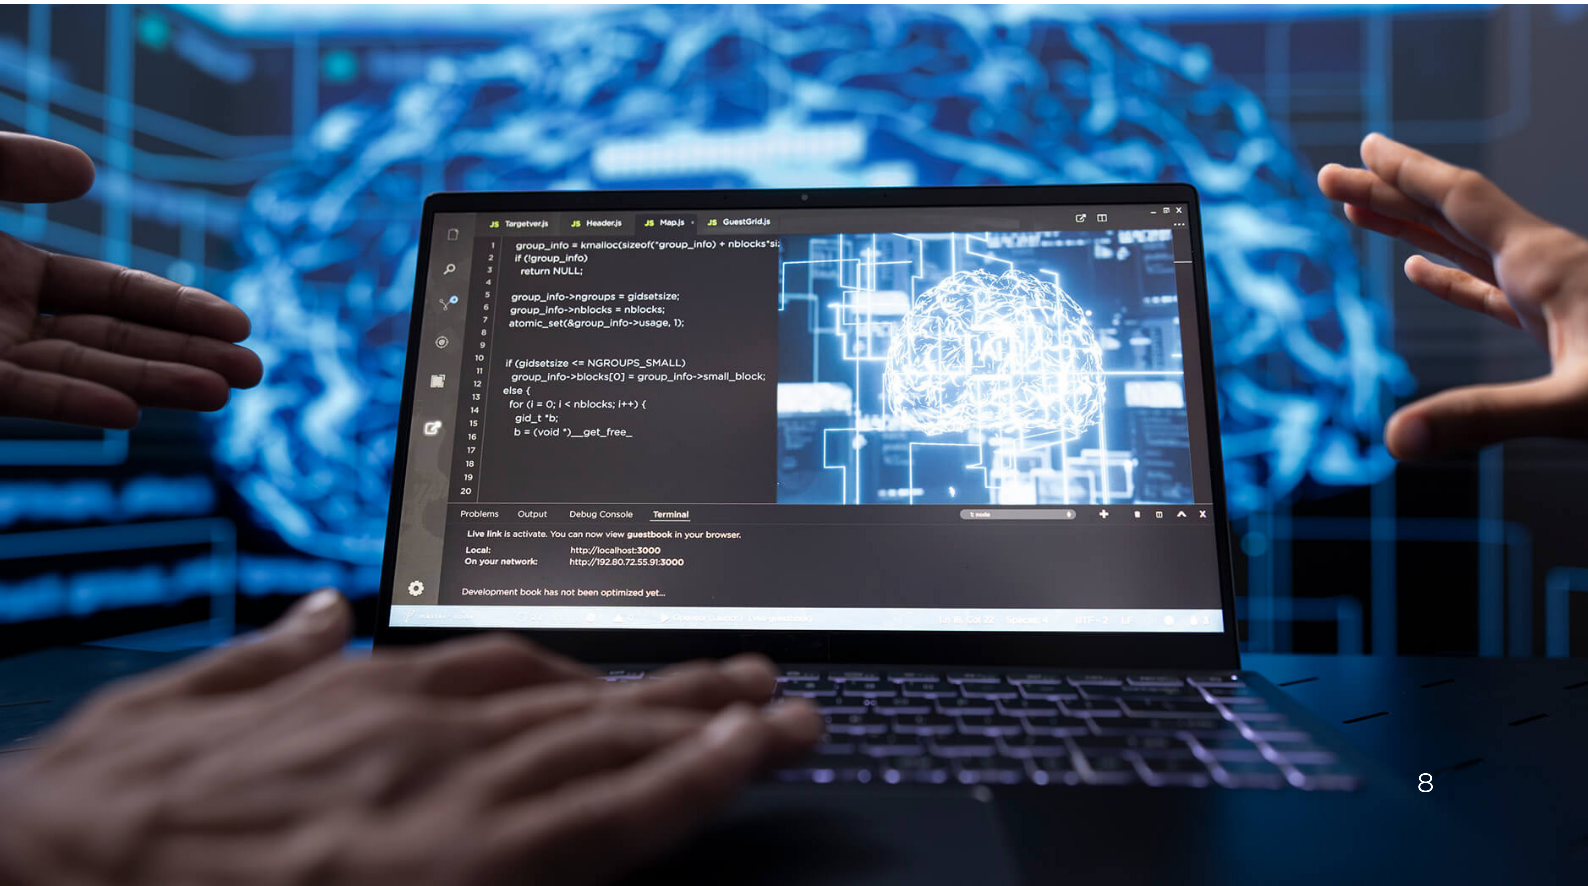

WRITER's 2026 Enterprise AI Adoption survey, conducted in partnership with Workplace Intelligence, sharpens the diagnosis through executive testimony rather than analyst projection (WRITER & Workplace Intelligence 2026). The study found that 79 percent of organizations face material challenges in scaling AI while 75 percent of executives admit, in their own words, that their company's AI strategy is more focused on appearance than execution. Nearly half (48 percent) describe their AI adoption to date as a massive disappointment. Together, these findings represent one of the clearest indications that AI investment and AI value have become decoupled at the executive level.

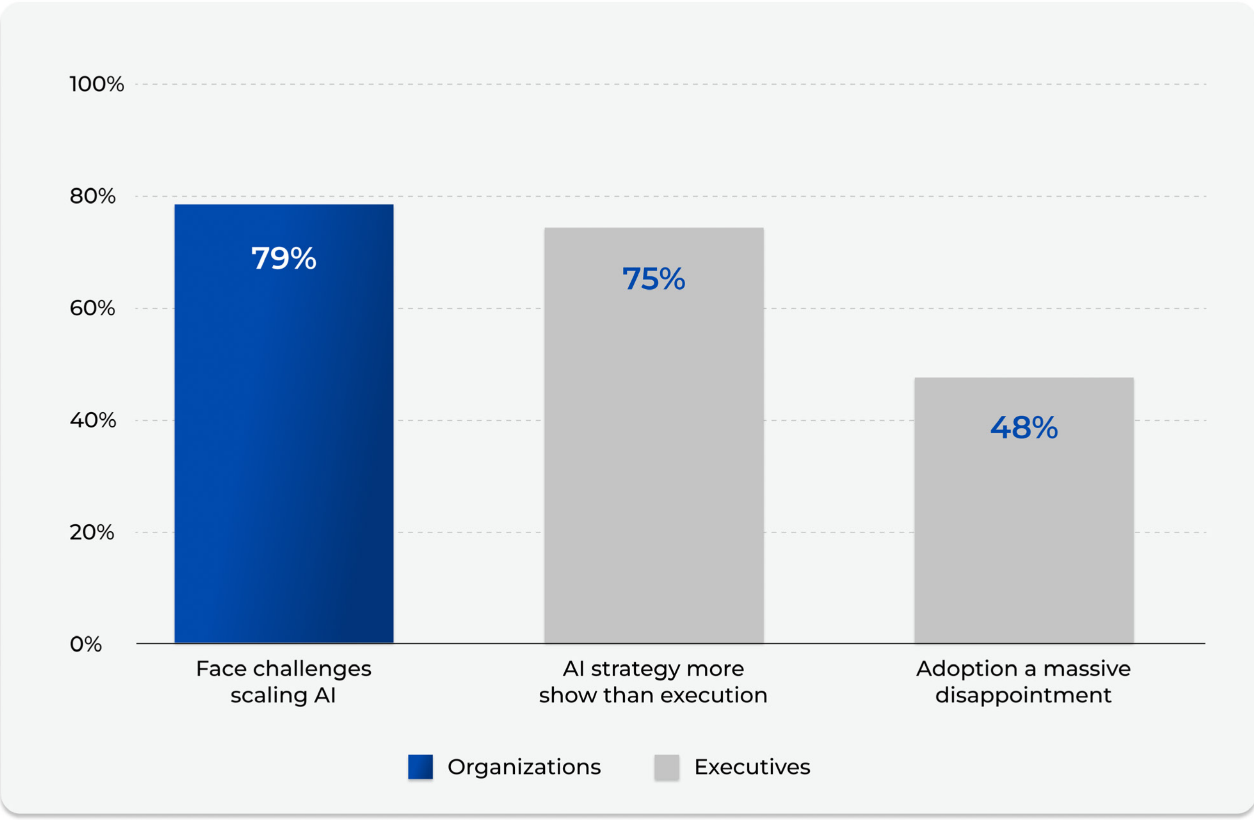

Figure 1. WRITER 2026 Enterprise AI Adoption survey findings by respondent group

Key Takeaway 1

Three quarters of executives, when asked privately, describe their company's AI strategy as more for show than for execution. This is not external criticism. It is leadership's own assessment of its own work, and it represents one of the clearest admissions to date that AI investment and AI value have decoupled.

Prosci, the change management research firm whose ADKAR model is used by approximately 80 percent of Fortune 100 companies, has surfaced a possible explanation for that disappointment (Prosci 2025). Nearly two-thirds of AI implementation challenges have nothing to do with the technology itself. The dominant failure modes are not algorithms, data quality, or computing capacity. Instead, Prosci's analysis points to people, training, trust, and the absence of structured change management as the primary causes of implementation failure.

Harvard Business Review reached the same conclusion in November 2025, arguing that most enterprises struggle to capture real value from AI because their people, processes, and politics fail to align around the technology (Li et al. 2025). Fear of replacement, rigid workflows, and entrenched power structures quietly derail AI initiatives, even in companies with advanced technology stacks. This pattern appears consistently across major research efforts conducted in 2025 and 2026. The bottleneck is cultural, not computational.

A different viewpoint deserves a place in any honest diagnosis because the cultural explanation does not account for every failure. Some deployments fall short because the workflow itself is poorly suited to current AI capabilities, and no amount of change management can redeem a tool that cannot reliably perform the underlying task. Starbucks provides a recent illustration of this dynamic (Reuters 2026).

In May 2026, the company retired an AI-powered inventory-counting tool built by NomadGo across more than eleven thousand company-operated stores in North America, nine months after deploying it under Chief Executive Brian Niccol as a centerpiece of the "Back to Starbucks" turnaround strategy. The vendor had promised 99 percent accuracy and claimed that it could count up to eight times faster than manual methods. The tool was intended to improve visibility into store-level shortages, yet it frequently miscounted and mislabeled items, confusing similar milk types or failing to recognize them altogether.

The deployment had executive ownership, a clearly defined outcome, and a measurable success metric in product availability. That ownership gave it the executive sponsorship the change management pillar requires, and the push for cleaner shortage data answered to the data foundation pillar. By those standards, it satisfied the business outcome discipline that anchors the framework's third pillar.

It fell short anyway because the computer vision challenge proved more difficult in a functioning coffeehouse environment than the vendor's claims suggested. Workers ended up checking nearly every count, so the tool doubled the task it was built to remove, and the net value turned negative. The deeper lesson is that the shortfall was catchable inside the framework. Section 8.5 carries the deployment discipline that's required across all three culture pillars: a staged rollout that tested the tool in the messiest real-world conditions and measured net value before scaling would have surfaced the miscounting in a handful of stores rather than across the entire North American fleet, which is exactly what section 8.5 makes explicit. Business outcome discipline set the right target; change management secured the sponsorship and data foundation framed the goal, and the deployment discipline in section 8.5 is what would have caught the mismatch between the workflow and the tool before it reached eleven thousand coffeehouses.

The MIT NANDA study reinforces the same caution at scale (MIT Project NANDA 2025). Drawing on 52 executive interviews, 153 leader surveys, and 300 public deployments, the research found that 95 percent of generative AI pilots delivered no measurable profit-and-loss impact. Only 5 percent of integrated systems generated significant value.

The lesson for leaders is twofold. The three pillars presented in the paper are necessary, and the evidence supports that claim. They are not, however, sufficient on their own. Workflow fit and an honest assessment of what current AI systems can realistically accomplish are prerequisites that the pillars assume rather than guarantee. The discipline to decline a deployment is just as important as the discipline to govern one. That principle applies as much to tasks AI cannot yet perform reliably as it does to initiatives that lack a named owner, a defined outcome, or a clear path to value.

2.

The 70% Rule

The 10/20/70 framework from Boston Consulting Group is the single most important diagnostic to emerge from the past two years of enterprise AI research, and it deserves a place in every board discussion of AI strategy (BCG 2025a). Algorithms account for 10 percent of the work involved in AI transformation. The technology backbone, including infrastructure, data pipelines, and model deployment, accounts for 20 percent. The remaining 70 percent comes from people and processes. That 70 percent is where enterprise AI strategies either compound or collapse, yet it is also where most enterprise AI budgets are systematically underfunded relative to the value at stake.

The framework describes where the work concentrates, not a fixed budget ratio every organization should hit. Some companies operating at the frontier intentionally carry significant infrastructure and model costs. For organizations with deeply empowered users and real adoption depth, that allocation may be entirely healthy. The misallocation BCG identifies is the more common scenario: investment accumulates in models and infrastructure precisely because the more difficult work of organizational change, adoption, and operating model has never begun.

Key Takeaway 2

Most enterprise AI budgets are inverted relative to where value compounds. The typical allocation pours capital into the 30 percent of algorithms and infrastructure while leaving the 70 percent of people and processes underfunded. The companies generating measurable AI value have rebalanced their budgets around this insight.

BCG's September 2025 Build for the Future report quantified the consequences of the misallocation in stark terms (BCG 2025b). Five percent of companies globally now qualify as future built for AI, meaning they are at the forefront of AI innovation, systematically building cutting-edge AI capabilities across functions, and consistently generating substantial value. Thirty-five percent are classified as scalers, meaning they are beginning to generate value but admitting they could move faster. The remaining 60 percent are classified as laggards, reporting minimal revenue and cost gains while lacking the capabilities required to scale AI effectively.

The gap between these three tiers is widening every quarter, and the dynamics of the next AI wave will widen it further rather than closing it. Generative AI produces content in response to prompts, with humans directing each step of the process. Agentic AI plans and executes multi-step tasks toward a defined objective, calling tools and making decisions with limited human intervention. The shift moves AI from a tool that people operate to a system that can increasingly operate on their behalf.

Agentic AI, the next layer of enterprise AI capability, already accounts for 17 percent of total AI value in 2025 and is expected to reach 29 percent by 2028, according to BCG's December 2025 analysis (BCG 2025a). The companies that have not yet built the cultural and data foundations required to absorb generative AI will be poorly positioned to capture value from agentic AI. The foundational requirements remain largely the same. This compounding effect is one of the most underappreciated dynamics in enterprise AI strategy and one of the single strongest arguments for investing in the three pillars now rather than waiting for the next platform shift.

Key Takeaway 3

Agentic AI will widen the gap, not close it. The conventional board narrative assumes laggards can catch up when the next AI wave arrives. The data points in the opposite direction because the cultural and data foundations required for agentic AI are the same ones required for generative AI. The catch-up window is narrowing, not widening.

The three-pillar framework introduced in this paper is the operational response to the BCG diagnostic. Each pillar addresses a specific failure mode in the 70 percent, and each is anchored by an enterprise that has executed the approach at scale. The goal is to provide leaders with real operating models rather than theoretical constructs.

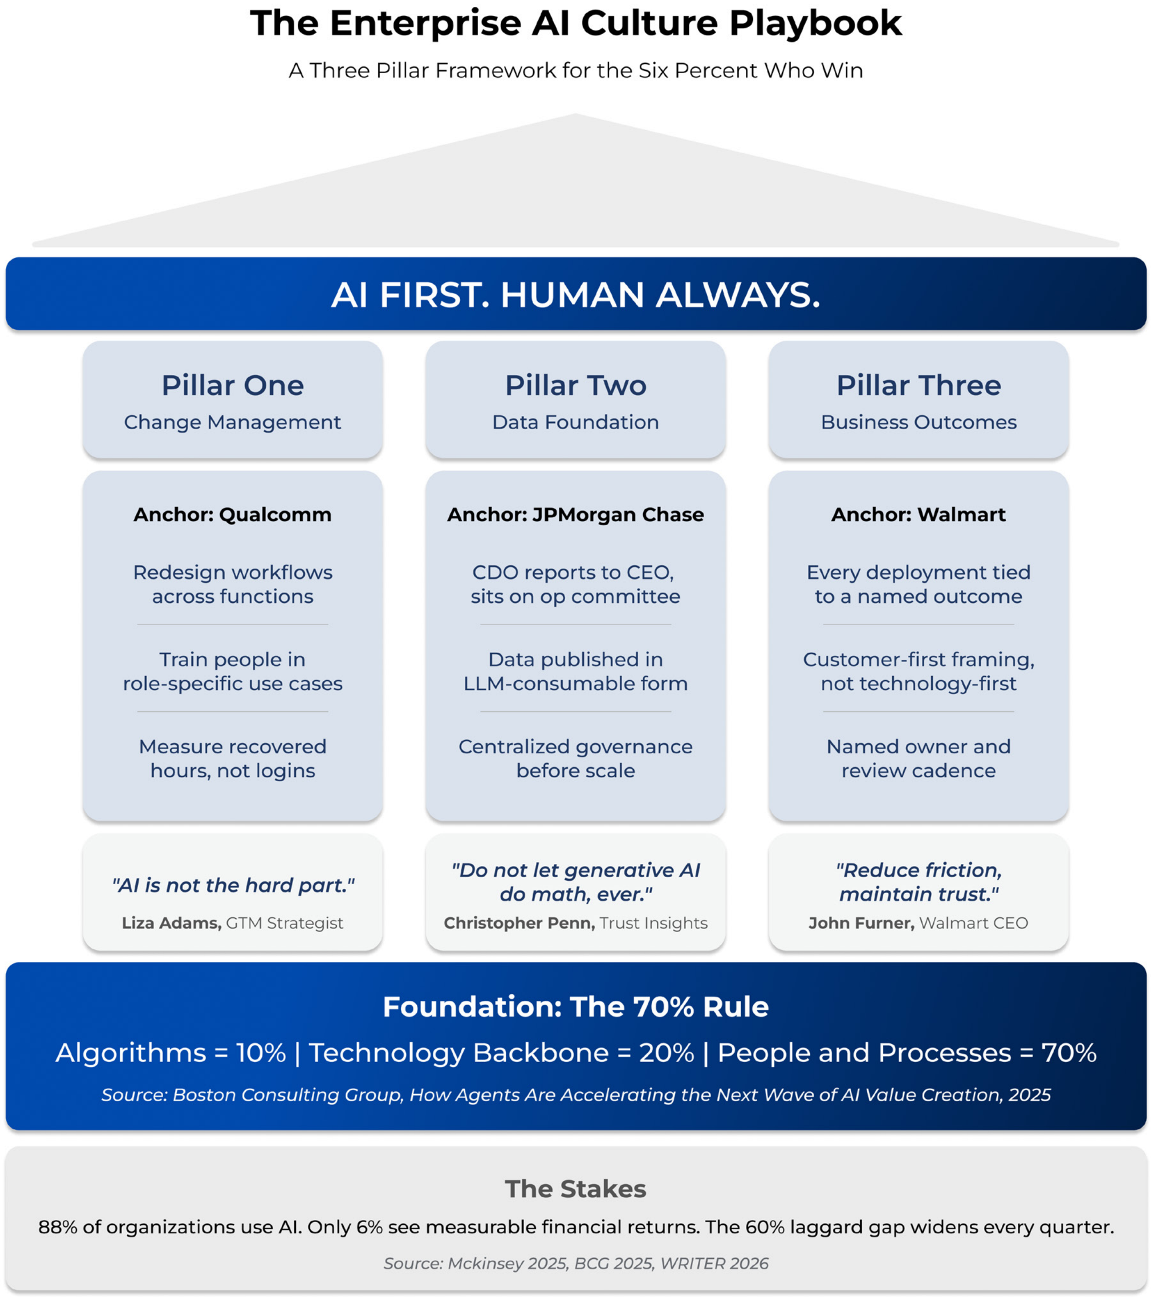

Figure 2. The Enterprise AI Culture Playbook framework, anchored by the BCG 70% Rule

"AI is not the hard part. People only build what they believe AI can do, and belief comes from seeing it work in their day-to-day. Doing the same work faster is the floor. The growth comes when people build what wasn't possible before. We can't reimagine the future by automating the past."

Liza Adams

3.

Pillar One: Change Management

3.1 Anchor Case: Qualcomm with WRITER

Most enterprises treat AI adoption as a waterfall software rollout, which is precisely why most enterprises struggle to extract meaningful value from their AI investments. The default playbook is familiar: procure the license, distribute the training email, measure user logins, and declare victory at the first sign of activity. None of those steps, however, fundamentally changes how work gets done. AI adoption that fails to redesign the underlying workflow cannot produce the productivity gains that justified the investment in the first place. Agile and lean methodologies, built around fast iteration and close attention to user needs, are far more likely to surface operating model gaps early while they remain inexpensive to correct.

Qualcomm took a different approach in partnership with WRITER, treating the rollout as a cultural redesign rather than a software installation (Qualcomm & WRITER 2025). They started with a pilot in marketing and communications, where participants expressed a desire to keep using the platform full-time. From there, the company scaled adoption through cohort-based training, office hours, and hands-on use-case workshops, the opposite of a one-time training email.

The semiconductor leader rolled out AI capabilities to hundreds of users across marketing, communications, legal, product, analytics, sales, learning and development, and human resources. Along the way, the organization vetted more than twenty-five distinct use cases and codified seventy unique workflows. Context and workflow design sat at the center of the effort, which is precisely the layer many enterprise leaders skip.

The result, as reported in the WRITER 2025 enterprise adoption case study, was approximately 2,400 hours of recovered productive time each month across participating users, or roughly 29,000 hours annually. Those gains materialized because the work itself was redesigned, not because more employees logged into a platform.

The Qualcomm case is instructive less for what it includes than for what it omits because the omissions reveal where the discipline actually lives. There is no hero algorithm, no proprietary model, and no claim of ten times productivity gains anywhere in the case study, which is unusual for a vendor-centered adoption narrative. What the case does include is the disciplined work of identifying use cases, redesigning workflows, training employees across functions, and measuring recovered hours. Together, those activities constitute change management executed as an operating model rather than as a slide in a steering committee deck.

3.2 Practitioner Insight

Liza Adams, founder of GrowthPath Partners and an AI advisor to companies including Klaviyo, Cox Automotive, and WP Engine, makes a similar point (Adams 2026). She advises leaders rolling out AI to commercial teams to stop pushing tools and start demonstrating the work itself.

3.3 Why This Pillar Matters

Prosci's 2025 research on AI implementation failure makes the structural case for treating change management as a board-level discipline rather than a training function (Prosci 2025). With nearly two-thirds of AI implementation challenges rooted in people-related issues, even the most advanced AI platform will be underutilized, mistrusted, or rejected if structured change management efforts do not run in parallel to the technology deployment. The Qualcomm playbook is replicable across industries and enterprise sizes, and the required investment is modest relative to the cost of enterprise AI licenses that sit underutilized on employees' desktops.

4.

Pillar Two: Data Foundation

4.1 Anchor Case: JPMorgan Chase

Every enterprise AI conversation eventually crashes into data. Most executive teams acknowledge its importance, assume their data foundation is adequate, and move on to more interesting questions about model selection. The 2026 Informatica data leader survey reveals why that assumption is dangerous, and the dangerous part is not what most leaders expect (Informatica 2026). Sixty-five percent of employees believe the data behind their AI systems is solid, yet 75 percent of data leaders say those same employees require serious upskilling in data literacy. Seventy-four percent say they also require AI literacy.

Key Takeaway 4

The data risk in enterprise AI is not simply poor-quality data. It is the confident trust in unproven data. Sixty-five percent of employees believe the data is solid while 75 percent of data leaders say those same employees lack the literacy to evaluate it. The gap between perceived data quality and actual data readiness is where enterprise AI projects die, and most executives are not aware that the gap exists in their own organization.

JPMorgan Chase operates on the opposite premise, and the contrast with most enterprises is instructive (JPMorgan Chase 2025). The firm runs more than 450 AI use cases in production and has stated its intention to expand that number to 1,000 by the end of 2026. More than 200,000 employees were onboarded to internal AI tools within an eight-month window. Those numbers are not achievable without disciplined data foundation work running underneath every use case. The structural commitments JPMorgan has made to enable that work are visible in how the firm has organized its data leadership at the executive level.

Mark Birkhead, firmwide chief data officer at JPMorgan Chase, has been explicit about this operating model in interviews and industry coverage (Markets Media 2025). The firm established a centralized Chief Data and Analytics Office in 2024, with all data initiatives now sitting under that single umbrella rather than scattered across business units. The chief data and analytics officer reports directly to Jamie Dimon and serves on the operating committee. One of the office's primary responsibilities is modernizing enterprise data so it can be published in a form that is consistent, governed, and consumable by large language models.

The same fit-for-purpose discipline shows up in what JPMorgan keeps away from large language models. In core financial calculations, identical questions must produce identical answers. Repeatability is non-negotiable, which means probabilistic models cannot serve as the sole decision-maker layer. Independent research is blunt about why: language models configured for deterministic output still swing in accuracy by as much as 15% across identical runs (Atil et al. 2024). The institutions extracting real value pair model capabilities with deterministic validation and reserve language models for the tasks they actually fit while the broader market pushes them into every workflow regardless (Daloopa 2026).

It is worth noting what sits alongside that structure. Dimon also brought tens of thousands of employees back to the office five days a week, defending it as an apprenticeship culture built on mentorship and collaboration. He has not framed that as an AI decision, yet proximity and culture are the raw material every change management effort runs on, and a firm this deliberate about the 70% is unlikely to see the two as unrelated.

Key Takeaway 5

The most consequential JPMorgan AI decision was a structural one. In most enterprises, data leadership is buried three layers down in IT. JPMorgan's chief data and analytics officer reports directly to the chief executive and serves on the operating committee, which is proof that data must sit at the executive table for AI to scale. The structure is replicable, and it is the single highest leverage move most boards have not yet made.

4.2 Why This Pillar Matters

The compounding cost of a weak data foundation is one of the most underappreciated risks in enterprise AI, and that risk compounds nonlinearly as AI deployment scales. Each additional AI use case deployed on top of brittle data multiplies the surface area for error, hallucination, and reputational exposure, which means the marginal cost of weak data rises faster than the marginal benefit of each new use case.

Enterprises that build the data foundation first, as JPMorgan has done, can scale use cases linearly while enterprises that deploy AI on top of an inadequate data foundation face exponentially compounding risk at every additional deployment.

5.

Pillar Three: Business Outcome Discipline

5.1 Anchor Case: Walmart

The third pillar separates real strategy from performance theater. For any enterprise AI initiative, the diagnostic question is direct: Is this deployment producing a measurable business outcome that the CFO can verify on the income statement, or is it producing an executive talking point that lives only in board materials?

In most enterprises, the honest answer is the latter. AI work that looks impressive in a press release but never reaches the financials is what executives now call theater, activity staged for an audience rather than results delivered to the business. The WRITER survey finding, with 75 percent of executives describing their own AI strategy in exactly those terms, landed with force for that reason.

Walmart's fiscal 2026 fourth-quarter earnings show what the alternative looks like (Walmart Inc. 2026). The company reported $713 billion in full-year revenue, up 4.7 percent year over year. Automation efforts in inventory management were credited as a material driver of sales performance, even as the company absorbed tariff-related headwinds through the year. The framing from Chief Executive John Furner provides the operational marker that distinguishes outcome discipline from performance theater (Furner 2026).

Furner did not lead his discussion of AI with the technology itself. He led with the customer experience that the technology enables. In his remarks to investors, Furner explained that the way Walmart is using technology and AI is helping the company create better customer solutions, reduce friction, simplify decision-making, and pinpoint inventory locations, all while maintaining the trust it has earned from customers and members.

Each of those outcomes is expressed in business terms rather than technological ones. Customer friction is reduced. Inventory accuracy is improved. Customer trust is preserved. Those are the measures of value Walmart is willing to be held accountable for.

Key Takeaway 6

The Walmart framing of AI as customer outcome infrastructure is the diagnostic test for outcome discipline. AI that cannot be tied to a named outcome, a named metric, a named owner, and a named review cadence should not be deployed. The Walmart case is replicable because the discipline is structural, not industry-specific.

5.2 Why This Pillar Matters

The WRITER 2026 survey finding that 75 percent of executives describe their AI strategy as more for show than for execution may be the most damaging admission in the current enterprise AI landscape, and the only durable corrective is outcome discipline executed at the deployment level (WRITER & Workplace Intelligence 2026).

Every AI deployment should be tied to a specific business outcome, with named ownership, named metrics, and a named review cadence. The Walmart framing of AI as customer outcome infrastructure provides a practical operating model.

Any deployment that cannot be tied to a measurable business outcome should be challenged before it is approved, and any deployment that cannot survive that scrutiny probably should not be deployed in the first place, and any deployment that cannot survive that scrutiny probably should not be deployed in the first place. The discipline of refusing to ship undisciplined deployments is the most leverageable behavior a chief executive can introduce.

6.

The AI Hollowing Risk

A distinct failure mode is emerging across enterprises that have deployed AI without first investing in the three pillars, and it deserves its own name because it is structurally different from healthy workforce evolution. The pattern is layoffs without redesign: cuts remove headcount while leaving the underlying workflows untouched. The result is a hollow organization that cannot absorb AI productively because the institutional knowledge required to direct AI was removed alongside the people who held it.

The data now confirms this is happening at scale. A February 2026 Careerminds survey of 600 HR professionals found that two in three employers who cut jobs for AI are already rehiring, often within months, and nearly a third reported losing critical skills and expertise in the process. Forrester puts employer regret at 55 percent. Only 8.4 percent of companies said their restructuring delivered on its promises.

This pattern is best described as AI Hollowing, a term I coined to capture what the numbers were revealing before they had a name. The term matters because it isolates the mechanism. The organization is not shrinking in a healthy, deliberate way; it is being hollowed out from the inside, with the load-bearing knowledge removed while the organizational shell remains.

Healthy workforce evolution follows a different sequence. Organizations invest in the three pillars first. They redesign workflows second. They redeploy talent into higher-value work third. The sequence matters because redesign and redeployment depend on institutional knowledge that is still present inside the organization. AI Hollowing reverses that order. It cuts first and hopes AI fills the gap afterward. The AI almost never fills that gap because the workflows it was expected to absorb were never redesigned to be absorbable in the first place.

Key Takeaway 7

AI Hollowing is a distinct failure mode that boards are not yet pricing into workforce decisions. Layoffs justified by projected AI productivity gains, without the workflow redesign occurring first, create organizations that cannot absorb AI. The projected savings often reverse because the institutional knowledge required to direct AI was removed alongside the people who held it.

Recent moves at Meta, Intuit, and Starbucks illustrate the variance in how enterprises are navigating this risk, and the market is reading the signals differently depending on whether the underlying culture and data work has been done.

Meta and Intuit announced roughly eleven thousand combined job reductions in 2026, though they framed them very differently: Meta tied its cuts to AI-driven efficiency gains and reorganized remaining employees into AI-focused units. Intuit's CEO, by contrast, attributed its 17 percent workforce reduction to organizational simplification and explicitly said that the decision had nothing to do with AI.

Starbucks, meanwhile, provides a different example. The company retired an AI inventory system deployed across more than eleven thousand stores because employees still had to verify every scan manually, effectively doubling the task rather than reducing it. At the same time, Starbucks continues to roll out other AI tools across the business

The diagnostic question for any enterprise considering AI-driven workforce reductions is whether the three pillars have already been built because workforce reductions are sometimes genuinely necessary. However, reductions made before work itself has been redesigned tend to remove capacity without removing the underlying workload. That is hollowing rather than productivity.

Done well, the sequence runs the other way: redesign the workflow first, confirm where capability is genuinely freed up, and then resize and redeploy talent into higher-value work. When that sequence is followed, the savings are real, and the downstream cost of AI hollowing is avoided.

7.

Compounding Economic Stakes

The widening gap between AI leaders and laggards has now reached the point where the difference is measurable in revenue growth, total shareholder return, and operating margin. The gap is no longer a forward-looking projection from analysts. It is showing up in reported financial performance.

PwC's 2026 AI predictions report found that visionary AI players are now reporting 1.7 times higher revenue growth, 3.6 times higher three-year total shareholder return, 2.7 times higher return on invested capital, and 1.6 times higher EBIT margin relative to laggards (PwC 2026). The performance gap has moved from forecast to financial fact, and it is visible on the balance sheet of companies that have built the three pillars, and those that have not.

Key Takeaway 8

The performance gap between AI leaders and laggards is already on the balance sheet. PwC's 2026 data show visionary AI players at 1.7 times higher revenue growth, 3.6 times higher three-year total shareholder return, and 1.6 times higher EBIT margin relative to laggards. This is a present financial fact, not a future forecast, and it changes how boards should evaluate AI investment in capital allocation decisions.

Independent research from Trust Insights, McKinsey, and Gartner confirms a consistent finding across the analyst landscape (McKinsey & Company 2025). Enterprises that move AI from pilots to production processes report an average return on investment of approximately 1.7 times, with cost savings in supply chain, finance, and customer operations now in the 26% to 31% range.

Forty percent of these enterprises expect positive yields within one to three years while another 35 percent expect returns within three to five years. which means three quarters of enterprises that build the three pillars and move to production should expect measurable financial returns within a five-year horizon.

The compounding dynamic is critical for board-level planning because the gap between the 5 percent of future-built companies and the 60 percent of laggards is not closing; it is widening every quarter (BCG 2025b). Every quarter spent optimizing the 10 percent of algorithms while underinvesting in the 70 percent of culture extends the gap further, and the window for catch-up investment is closing faster than most enterprise strategic plans assume.

The strategic implication for boards is direct: the cost of inaction now exceeds the cost of structured investment in the three pillars. The longer the inaction continues, the larger the gap to close becomes, and the more expensive it is to close.

8.

The Path Forward for Executive Leadership

The recommendations that follow are oriented to the C-suite and the board. They are deliberately structural rather than tactical because the tactical work belongs to functional leaders. The structural work belongs to the chief executive and the board, and the five moves below represent the structural decisions that distinguish the 5 percent of future-built enterprises from the 60 percent of laggards.

8.1 Restructure for the Three Pillars

Each of the three pillars requires named executive ownership at the operating committee level. That ownership cannot be delegated to functional leaders without losing the structural leverage that makes the framework work.

Change management should sit with a senior leader empowered to redesign workflows across functions rather than with a training function inside human resources, where the authority to redesign work often does not exist. Data foundation should sit with a chief data and analytics officer reporting directly to the chief executive, following the JPMorgan model. Business outcome discipline should sit with the chief operating officer or equivalent, with explicit accountability for tying every AI deployment to a named metric.

8.2 Audit the 70%

The diagnostic exercise most enterprises have not yet undertaken is a structured audit of where current AI investment is being allocated against the BCG 10/20/70 framework, and the audit can usually be completed within a single quarter (BCG 2025a).

The expected finding in most enterprises is dramatic underinvestment in the 70 percent, with the bulk of spend concentrated in models, infrastructure, and vendor contracts rather than in change management, data foundation, and outcome discipline. The corrective is to rebalance the next budget cycle around people, processes, and operating models rather than around model selection and infrastructure procurement, which is a board-level decision rather than a functional one.

8.3 Tie Every Deployment to a Business Outcome

No AI deployment should proceed without the following elements in place. The discipline of refusing to ship undisciplined deployments is one of the most leverageable behaviors a chief executive can introduce into the AI program:

- A named business outcome
- A named metric
- A named owner
- A named review cadence

The Walmart framing of AI as customer outcome infrastructure is the operational model, and the model is industry agnostic (Walmart Inc. 2026). Any deployment that cannot pass these four criteria is unlikely to deliver measurable value and should be deferred until the criteria can be met because the cost of shipping undisciplined deployments compounds across the portfolio as quickly as the cost of a weak data foundation compounds across use cases.

These four criteria decide whether a deployment is worth pursuing. Section 8.5 governs how to roll it out without betting the enterprise on an unproven tool.

8.4 Avoid AI Hollowing

Workforce reductions tied to AI productivity claims should be paused until the three pillars are in place because the downstream cost of hollowing materially exceeds the upfront savings of the reduction. Gartner's 2026 research points the same way, finding that the firms capturing the most AI value were not necessarily the ones cutting the most. The strongest returns came from amplifying workers rather than replacing them.

The structural response is to make AI-driven workforce reductions conditional on a documented audit of the three pillars, with the chief executive personally accountable for the audit and the board reviewing the audit before approving the reduction.

8.5 Stage the Rollout and Scale on Proof

The three pillars define the foundations an enterprise builds before it deploys. They do not, by themselves, govern the act of deploying, and the Starbucks inventory case shows what that gap costs. That deployment had executive ownership, a named outcome, and a clear metric, which means it would have passed the five criteria in section 8.3. Yet it scaled to more than eleven thousand stores before anyone confirmed the tool produced net value in a working coffeehouse environment. The missing discipline was staged validation. The pilot threshold was never set high enough to catch a tool that doubled the counting task rather than removing it.

Deployment discipline closes that gap, and it runs across all three pillars rather than replacing any of them. Five practices operationalize it:

- Stage every rollout through defined gates, piloting in a representative sample of the messiest real-world conditions before any fleet-wide scale.
- Set the success bar and the stop criteria before the pilot begins, and define value as net of the human cost the tool creates, including verification and rework.
- Test under worst-case operating conditions rather than vendor demonstration conditions because the distance between the two is where deployments fail.
- Wire frontline feedback into the go or no-go decision since the people doing the work see the failure first.
- Time-box each gate. A multi-quarter, full-scale deployment that is later retired is evidence that the gate was never truly there.

Evidence earns the rollout. A deployment proves itself by clearing the pilot bar under real operating conditions, and the discipline of holding the line at the gate is what converts a hard-task failure from a fleet-wide write-off into an inexpensive lesson learned in a handful of stores.

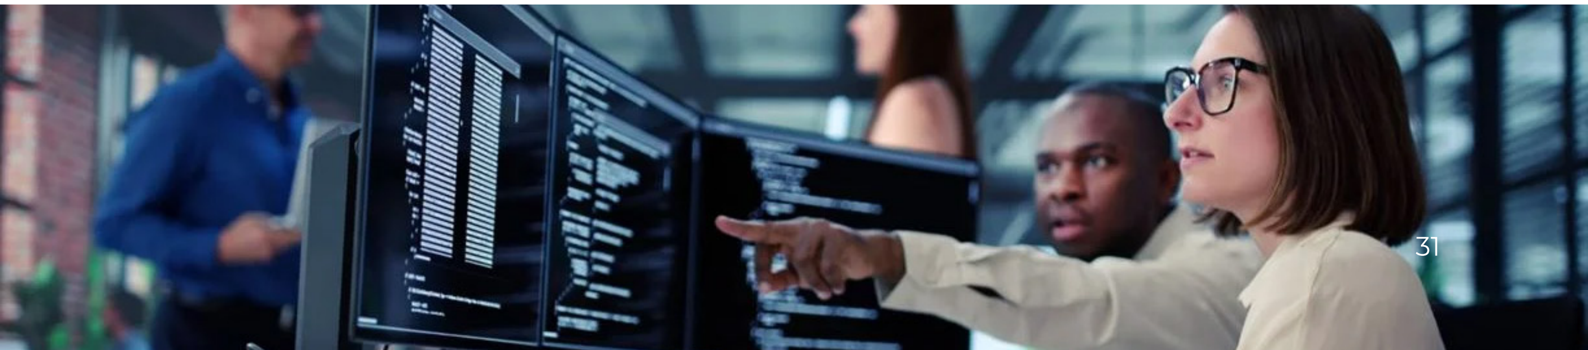

9.

A Leadership Charge

The question facing every enterprise executive in 2026 is not whether to deploy AI because that question was answered two years ago by competitive pressure and capital markets. The actual question is whether the cultural, data, and outcome discipline of the enterprise is ready to absorb AI at scale and convert it into compounding value over a multi-year horizon. The companies that answer affirmatively will continue to widen the gap while the companies that defer the answer will continue to fall behind on the measurable financial metrics that boards and investors are now tracking.

Linda Du, founder and chief executive of Moola Money and managing director of Okta Investment in Berlin, framed the leadership posture required for this moment with unusual clarity in her Thrive Global conversation (Du 2025). Du argues that leadership starts with self-discipline and the ability to set a vision and execute it step-by-step. Leaders should not wait for permission. If there is an idea or a vision present, the right move is to start now rather than wait for the conditions to feel safer. European founders building AI ventures from scratch are operating on that clock, and global enterprises should be operating on the same clock because the window for repositioning around the three pillars is narrower than most strategic plans assume.

The technology is ready, the infrastructure is ready, and the models are ready, which means none of the three is the constraint on enterprise AI value creation in 2026. The remaining question, the one that will determine which enterprises emerge from this decade as future built and which will be reclassified as laggards, is whether the culture is ready to absorb what the technology now makes possible. The framework introduced in this paper is the operational answer to that question, and the five structural recommendations are the moves a chief executive can make this quarter to begin closing the gap before the gap becomes uncatchable.

AI first, human always. AI creates the possibility; people create the value.

References

1. Adams, L. 2026. “AI Transformation Begins with People: Reimagining Work with Liza Adams.” The CMO. https://thecmo.com/career/ai-transformation-begins-with-people/.
2. Berk Atil et al. “Non-Determinism of ‘Deterministic’ LLM Settings.” arXiv:2408.04667. https://arxiv.org/abs/2408.04667.
3. Boston Consulting Group. 2025a. “How Agents Are Accelerating the Next Wave of AI Value Creation.” BCG. https://www.bcg.com/publications/2025/agents-accelerate-next-wave-of-ai-value-creation.
4. Boston Consulting Group. 2025b. “The Widening AI Value Gap: Build for the Future 2025.” BCG. https://www.bcg.com/publications/2025/are-you-generating-value-from-ai-the-widening-gap.
5. Careerminds. “AI Layoffs and Rehiring: A Survey of 600 HR Professionals.” https://careerminds.com/blog/cost-of-ai-layoffs.
6. Carter, S. 2025. *AI First, Human Always*. Wiley. https://aifirstbiz.com/.
7. Daloopa. How to Use LLMs for Financial Data Analysis: Complete Implementation Guide. https://daloopa.com/blog/analyst-best-practices/exploratory-financial-data-analysis-using-large-language-models.
8. Du, L. 2025. “Mindful Leadership, Inside Linda Du’s Double Venture. *Thrive Global*. https://community.thriveglobal.com/leading-with-heart-the-power-of-mindful-leadership.
9. Furner, J. 2026. “Walmart Q4 FY2026 Earnings Call.” Walmart Inc. https://www.linkedin.com/posts/sellcord_walmart-q4-fy2026-report-is-live-activity-7432493770346205185-D8Ke/.
10. Gartner. “The Hidden Workforce Costs of AI.” https://www.gartner.com/en/articles/ai-workforce-costs.
11. Informatica. 2026. “2026 CDO Insights Report: AI Adoption Trends, Trust, Data Quality, and Governance Challenges.” Informatica. https://www.informatica.com/lp/cdo-insights-2026_5264.html?formid=43815.
12. JPMorgan Chase. (2025). *Annual Report: Form ARS FY2025*. JPMorgan Chase & Co. https://jpmorganchaseco.gcs-web.com/node/872291/pdf.

13. Li, J., Zhu, F., and Hua, P. 2025. “Overcoming the Organizational Barriers to AI Adoption.” Harvard Business Review. https://hbr.org/2025/11/overcoming-the-organizational-barriers-to-ai-adoption.
14. Markets Media. 2025. “JPMorganChase Makes Data AI-Ready. Markets Media. https://www.marketsmedia.com/jpmorganchase-makes-data-ai-ready/.
15. McKinsey & Company. 2025. “The State of AI 2025.” McKinsey & Company. https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai.
16. MIT Project NANDA. “The GenAI Divide: State of AI in Business 2025.” https://nanda.media.mit.edu/.
17. Penn, C. S. 2025. “How to Transform Data into Actionable Intelligence.” MAICON 2025 Speaker Series, Trust Insights. https://www.marketingaiinstitute.com/blog/marketing-ai-conference-christopher-penn.
18. Prosci. 2025. “Best Practices in Change Management: 2025 Edition.” Prosci. https://empower.prosci.com/best-practices-change-management-executive-summary.
19. PwC. 2026. “2026 AI Business Predictions.” PricewaterhouseCoopers. https://www.pwc.com/us/en/tech-effect/ai-analytics/ai-predictions.html.
20. Qualcomm & WRITER. 2025. “Enterprise AI Adoption Case Study. WRITER. https://writer.com/blog/enterprise-ai-adoption-2026/.
21. Reuters. “Starbucks Retires AI Inventory Counting Tool Across North American Stores.” https://www.reuters.com/business/starbucks-scraps-ai-inventory-tool-across-north-america-2026-05-21/.
22. Walmart Inc. 2026. “*Form 8-K: Q4 FY2026 Earnings Release*.” US Securities and Exchange Commission. https://stock.walmart.com/_assets/_461d6b46a29d437b51015f942ff9bb4e/walmart/db/938/9972/earnings_release/Earnings+Release+(FY26+Q4).pdf.
23. WRITER & Workplace Intelligence. 2026. “2026 Enterprise AI Adoption Survey: Why 79% Face Challenges Despite High Investment.” WRITER. https://writer.com/blog/enterprise-ai-adoption-2026/.

Author

Sandy Carter

Sandy Carter serves as the Non-Executive Chair of the Applied AI workgroup at The Digital Economist's Center of Excellence on Human-centered Global Economy. She is the bestselling author of AI First, Human Always, and a Forbes contributor covering AI, agentic commerce, blockchain, and enterprise technology. She has built multi-billion dollar businesses at IBM and Amazon Web Services, served on the board of Altair Engineering through its $11 billion acquisition by Siemens, and has spoken in ninety-two countries at venues including Davos, the World Bank, and the United Nations. Sandy founded Unstoppable Women of AI and Blockchain, a community that has trained more than fifty-five thousand women across ninety-two countries.

About

The Digital Economist, headquartered in Washington, D.C. with offices at One World Trade Center in New York City, is the world's foremost think tank on innovation advancing a human-centered global economy through technology, policy, and systems change. We are an ecosystem of 40,000+ executives and senior leaders dedicated to creating the future we want to see—where digital technologies serve humanity and life.

We work closely with governments and multi-stakeholder organizations to change the game: how we create and measure value. With a clear focus on high-impact projects, we serve as partners of key global players in co-building the future through scientific research, strategic advisory, and venture build out.

We engage a global network to drive transformation across climate, finance, governance, and global development. Our practice areas include applied AI, sustainability, blockchain and digital assets, policy, governance, and healthcare. Publishing 75+ in-depth research papers annually, we operate at the intersection of emerging technologies, policy, and economic systems—supported by an up-and-coming venture studio focused on applying scientific research to today's most pressing socio-economic challenges.

CONTACT: INFO@THEDIGITALECONOMIST.COM

CENTER OF EXCELLENCE

The Digital Economist Center of Excellence for a Human-Centered Global Economy is dedicated to addressing the biggest challenges humankind and our planet face by leveraging digital technologies for good.

The Digital Economist Executive Fellowship invites senior leaders and decision-makers to join our Center of Excellence, providing them with a platform for amplification and global impact. This unprecedented, one-of-a-kind opportunity enables Executive Fellows to network and build relationships at the highest level, driving transformative change and innovation in the global digital economy.

The Executive Fellowship

The Digital Economist Executive Fellowship is a selective leadership program integrating visionary professionals into the Center of Excellence for a Human-Centered Global Economy to advance global economic policy and systems transformation.

Global Impact

Amplify your influence and drive transformative change by participating in high-level initiatives that address the most pressing global challenges.

Elite Community

Become part of an exclusive network of visionary leaders and innovators, collaborating to shape the future and drive global progress.

Unparalleled Opportunities

Access unique platforms and events that enhance your professional journey, providing unparalleled opportunities for growth, visibility, and leadership.

Participation Framework

Time Commitment

Minimum commitment: 24 hours per year, for the monthly Center of Excellence meetings. On-demand consultation with the Fellowships team.

Publications

Executive Fellows are expected to contribute to two key publications per year, launched at key global events such as Davos and New York Climate Week.

In-Person Convenings

Executive Fellows are invited to in-person convenings in North America and Europe, with regional convenings in Africa, Latin America, and Asia.

Speaking Engagements

Executive Fellows are offered speaking opportunities throughout the year to amplify their work and contributions.

Our Executive Fellows are at the forefront of research, policy discourse, and systems-level transformation.

- Applied Artificial Intelligence
- Digital Assets & Blockchain
- Sustainability in Tech
- Tech Policy & Governance
- Quantum Computing
- Cyber Studio
- Regenerative Digital Infrastructure
- Healthcare Innovation

Publications

Ideas that shape the future.

The Digital Economist's publications translate research into high-signal outputs: frameworks, policy papers, and industry outlooks that advance a sustainable, inclusive digital economy and inform decision-making across markets and institutions.

Explore our full portfolio of publications and research outputs:
www.thedigitaleconomist.com/publications

Engagement Opportunities

Executive Fellows have access to over **500 events globally** in a Fellowship cycle.

Executive Leadership Dinner London Climate Action Week
June 24, 2026

The Digital Economist Executive Retreat
August 21–23, 2026

United Nations General Assembly (UNGA 81)
September 8–22, 2026

Climate Week NYC 2026
September 20–27, 2026

World Bank Group and IMF Annual Meetings
October 12–18, 2026

Join the Fellowship

Advance your leadership within a global platform shaping technology, policy, and economic systems transformation.

Access Full Brochure | **Apply Now** | **Learn More**

Institutional Research Network

A Fragmented World Requires New Institutional Leadership

Technology, economics, and governance are shifting faster than traditional institutions can adapt. AI ecosystems, digital assets, geopolitical competition, sustainability transitions, and new governance architectures demand clarity, legitimacy, and a coherent strategy.

Institutions must now operate as signal generators—shaping the narratives, norms, and systems that define global markets.

Why We Built the Institutional Research Network

A global research and convening platform enabling institutions to:

- Shape emerging policy and governance discourse
- Build narrative power in a volatile environment
- Co-author high-signal research with global experts
- Gain visibility at the world's most influential convenings
- Anchor strategy in human-centered, future-forward frameworks

Co-Authorship & Knowledge Pathways

Through structured co-authorship across eight priority domains—Tech Policy and Governance, Digital Assets & Blockchain, Sustainability in Tech, Applied Artificial Intelligence, Cyber Studio, Quantum Computing, Regenerative Digital Infrastructure, and Healthcare Innovation—institutions contribute to high-level research that informs policy dialogue, regulatory development, and strategic decision-making.

Participation extends beyond commentary. Institutions are integrated into published research, roundtable dialogues, and domain-specific working groups that inform regulatory discussions and industry standards. This structured engagement enables organizations to contribute at the research and drafting stage, engage directly with policymakers and industry leaders, and align internal strategy with emerging policy and market developments, resulting in active presence within decision-making environments rather than passive visibility.

We invite your organization to schedule a strategic briefing to map research priorities and determine the appropriate integration pathway within the Institutional Research Network.

Reach us at partnerships@thedigitaleconomist.com.
Visit us at **thedigitaleconomist.com**

The Digital Economist Ventures

Applied Platforms. Strategic Domains. Real-World Implementation.

Research defines the questions. Ventures test the answers.

In addition to research and convening, The Digital Economist advances a portfolio of venture platforms that extend inquiry into applied domains, where governance, infrastructure, and market design move from dialogue to deployment.

Each venture operates with a defined mandate while remaining integrated within the broader institutional ecosystem.

Tech for Transparency

Financial integrity in the digital age

Advances financial accountability and anti-corruption frameworks through distributed technologies and data-driven transparency systems. Positioned at the intersection of blockchain infrastructure and institutional reform, it translates transparency principles into operational tools.

The Ostrom Project

Reimagining digital commons governance

Explores collective stewardship models for emerging digital systems. Drawing on principles of shared resource governance, it develops frameworks for sustainable digital infrastructure and cooperative system design.

ANER-G

Energy systems innovation

Focuses on decentralized infrastructure, programmable energy markets, and next-generation grid integration. It addresses the structural evolution of energy systems within digital and blockchain-enabled environments.

Africa Coalition

Continental coordination for strategic sectors

Convening leaders across energy, infrastructure, finance, health innovation, education, and future capabilities, the Coalition creates structured engagement pathways for continental collaboration.

Explore the full ecosystem at thedigitaleconomist.com

THE
DIGITAL
ECONOMIST