Hype Is Not a Strategy

A Framework for AI Adoption That Actually Works

How enterprise leaders can move from speculative AI spend to structured, measurable, and defensible transformation.

EXECUTIVE SUMMARY

> **The problem isn't AI. It's the absence of structure around it.**

Something is happening inside every enterprise right now. Employees are using AI tools you didn't deploy. Budgets are getting absorbed by initiatives that never reached production. Productivity is supposed to be up, but the delivery numbers aren't moving. And when the board asks for a report on AI program performance, no one has a clean answer.

This isn't a technology problem. There are people making AI work. The honest truth behind the missed outcomes and waste is that most organizations adopted AI without the structural foundation to make it succeed. No defined goals. No baseline to measure against. No way to know which tools employees are actually using, or whether the training they received prepared them to use those tools well.

Hype drove the investment. Urgency drove the rollout. And now a significant share of that spend is sitting in a black box, generating activity but not outcomes.

This paper describes a framework for doing it differently. Not slower — differently. It covers how outcome measurement, workforce readiness, and automated guardrails work together as a continuous cycle rather than three separate initiatives. And it gives

leaders a practical entry point for moving from AI adoption anxiety to structured, defensible transformation.

> **We move clients from AI Adoption Anxiety to the relief of structured outcomes — helping to create deep interconnected foundations that ensure organizations don't break under pressure.**

PART 1

Carried, Not Leading

The dominant story about enterprise AI is one of momentum — the sense that organizations must keep moving or be left behind. And that story is doing real damage.

Fear of falling behind is not a strategy. It generates the kind of adoption that produces activity without accountability: tools deployed without governance, budgets approved outside standard procurement, training that doesn't connect to the tools employees are actually using. The rush and momentum are real. The destination is not.

Three symptoms have become consistent across organizations that adopted AI under pressure rather than by design.

Symptom 1: The Returns Aren't Moving

AI spend is growing. The returns, for most organizations, are not. The median ROI from enterprise AI initiatives currently sits at just 10% — well below the 20% that most finance leaders had originally planned for. 56% of CEOs report that AI adoption has produced neither higher revenues nor lower costs over the last twelve months. And when organizations look for a meaningful impact at the enterprise level, only 39% find one.

The deeper problem: most organizations never set a number before they started. 60% of companies have no defined target for AI impact. Without a baseline, there's nothing to measure against — and no legitimate way to know whether an initiative is working, stalled, or quietly wasting budget.

45%
of orgs can quantify AI ROI
BCG, 2025

60%
have no defined AI impact target
MCKINSEY, 2025

46%
of AI projects scrapped before adoption
S&P GLOBAL, 2025

56%
of CEOs: no revenue or cost benefit from AI
PWC, 2026

Symptom 2: The Workforce Isn't Ready

Only 12% of employees report receiving sufficient training to unlock meaningful productivity gains from AI. The rest are using tools they weren't trained for, producing outputs they can't validate, and absorbing the rework when those outputs fail.

The consequences are not hypothetical. A 2025 study found that developers using AI tools were 19% slower at real-world tasks — while believing they were 20% faster. A separate Anthropic study found that AI-assisted engineers scored 17% lower on comprehension and mastery assessments than those who coded without assistance.

64% of employees report that their workloads have actually increased since AI adoption. The productivity gain that was promised is, for most organizations, still theoretical. 54% of companies offer no AI training at all — organizations that skip readiness don't just underperform, they hollow out the workforce they need to recover.

12%	19%	64%	54%
of employees receive sufficient AI training	slower on real tasks while believing 20% faster	say workloads increased since AI adoption	of companies offer no AI training at all
EY, 2025	METR, 2025	EY, 2025	GENERAL ASSEMBLY, 2026

Symptom 3: The Risk Is Accumulating Out of Sight

65% of employees are already using AI tools without any centralized organizational approach. AI adoption across the enterprise sits at 95%. Formal guardrails coverage sits at 53% — and has grown by only 3% in the past year.

The gap between adoption and governance is not closing. It's widening. And the organizations most exposed to it are the ones whose employees are most engaged with AI — the ones using the most tools, running the most experiments, producing the

most output without oversight. Shadow AI doesn't eliminate risk. It just accumulates it out of sight, in the tools no one approved, the datasets no one audited, and the outputs no one validated.

72% of S&P 500 companies now flag AI as a material risk in public disclosures — up from just 12% in 2023. The board conversation has started. Most organizations don't yet have a clean answer.

65%	**95%**	**72%**	**42%**
use AI without centralized oversight	AI adoption rate vs. 53% formal guardrails	of S&P 500 flag AI as material risk	of orgs abandoned most AI initiatives YoY, up from 17%
AMA, 2025	AMA, 2025	CONFERENCE BOARD, 2025	S&P GLOBAL, 2025

PART 2

One Root Cause, Three Symptoms

The three symptoms above — absent ROI, unprepared workforce, accumulating risk — are not separate problems. They share a single root cause: organizations adopted AI without the structural foundation to govern, enable, or measure it.

The momentum was real. The urgency was real. But urgency and structure are not the same thing. Urgency drives speed of deployment. Structure determines whether that deployment produces outcomes.

Most enterprises treated AI as a category of tools to be acquired rather than a capability to be built. They procured before they governed. They deployed before they trained. They spent before they defined what success would look like. And now they are managing the downstream consequences of all three.

> **Hype is not a strategy. And urgency, without structure, is just expensive speed.**

The organizations that are generating real returns from AI — McKinsey identifies only 6% as genuine high performers — are distinguished by a single consistent practice: they moved from ad-hoc experimentation toward centralized, structured oversight. They are 3.6 times more likely to use AI for transformative change rather than incremental gains. They define outcomes before they spend. They train before they deploy. They govern proactively rather than reactively.

The gap between these organizations and the rest is not technical. It's structural. The question for leaders isn't whether to adopt AI — that decision has already been made by their employees, their competitors, and their boards. The question is whether to lead that adoption or be carried by it.

A structured approach to AI adoption rests on three disciplines, each addressing one of the symptoms above. They are not sequential. They are interdependent — and when they operate together as a continuous cycle, the compound effect is a program that learns, improves, and generates defensible value over time.

PART 3

A Framework for Structured AI Adoption

The three disciplines below are not a product roadmap. They are a framework for how disciplined AI adoption works at the organizational level — what it requires, what it prevents, and what it produces.

01 Measure First

OUTCOMES BEFORE SPEND

The discipline: define success before you start

The single most consistent predictor of AI ROI is whether an organization defined its baseline and target outcome before an initiative began. Not during. Not after. Before.

This sounds obvious. It rarely happens. 60% of organizations have no defined target for AI impact. 25% of organizations don't measure the value of AI at all. When an initiative fails to deliver, there is often no shared definition of what "delivering" would have looked like — which means nothing is learned, nothing is corrected, and the budget cycle repeats the same pattern.

The counterpoint is equally clear. AI leaders who embed revenue-focused ROI discipline at the proposal stage generate 2.1× greater ROI than their peers. Organizations with formal AI strategy and program visibility report an 80% implementation success rate, compared to 37% for those operating in silos. The discipline of measuring first is not bureaucracy. It's what separates programs that scale from programs that stall.

WHAT THIS LOOKS LIKE IN PRACTICE

- Every AI initiative defines a measurable baseline and a target KPI before a dollar of budget is committed
- A proposal review process surfaces competing initiatives and allows leadership to allocate resources to the highest-value opportunities
- Tracking is ongoing, not periodic — so organizations can scale what's working and stop what isn't before it consumes another quarter of budget
- A unified dashboard gives leadership a single, real-time view of every AI initiative, tool, and outcome across the organization

2.1× greater ROI for AI leaders who track financial KPIs at the initiative level. (BCG, 2025)

80% implementation success rate for organizations with centralized AI program visibility, vs. 37% in silos. (Writer, 2025)

62% of leaders admit to intentionally linking traditional software purchases to "AI initiatives" to bypass procurement controls. (Emburse, 2025)

---

02 Readiness Before Deployment

CAPABILITY BEFORE TOOLS

The discipline: enable the people before you deploy the platform

AI tools are only as effective as the people using them. This is not a controversial statement. It is simply ignored more often than it should be.

The most dangerous pattern in enterprise AI is what might be called the awareness gap: employees receive a login, a brief orientation, and the implicit expectation that they will figure out the rest. Most do figure out something — but what they figure out is

often not what the organization needs. They develop habits around the path of least resistance, not the path of highest value. They produce outputs they can't validate. They trust results they shouldn't.

The METR study finding that developers were 19% slower while believing they were 20% faster is not an anomaly. It is a measurement of the awareness gap in action. The Anthropic study finding a 17% drop in coding comprehension among AI-assisted engineers is the same gap, measured differently. Both point to the same structural failure: the tool was deployed before the capability was built.

The fix is not more training in the generic sense. Generic AI literacy courses produce awareness, not capability. The discipline that works is role-based enablement tied directly to the specific authorized tools an organization is running — so that employees learn to use the tools they actually have, for the workflows they actually own, in ways that are measurably safe and effective.

WHAT THIS LOOKS LIKE IN PRACTICE

- Training is linked to the specific AI tools an organization has approved — not generic AI awareness
- Role-based tracks ensure relevance: a developer's readiness path looks different from a finance analyst's
- Certification closes the loop — employees aren't just trained, they're verified as capable of safe, effective usage
- Readiness is tracked as an organizational metric, not a compliance checkbox

2× more likely to see significant positive AI ROI for organizations with mature upskilling programs. (DataCamp/YouGov, 2026)

40% of potential AI productivity gains are lost to inadequate workforce readiness. (EY, 2025)

14 hrs of weekly productivity gain for employees with 81+ hours of annual AI training, vs. 8 hours for those with minimal training. (EY, 2025)

---

03 The Fast, Safe Yes

GUARDRAILS THAT ENABLE, NOT BLOCK

The discipline: replace the slow no with a structured yes

The conventional response to AI risk is restriction. Block the unapproved tools. Require manual review for every request. Build a policy and enforce it through friction. This approach has one consistent outcome: it drives adoption underground.

People won't stop using AI. They simply hide it. When the approval process takes weeks and the alternative is a personal subscription no one can see, the choice is not difficult for most employees. The result is Shadow AI — 65% of employees using tools without organizational oversight, creating data exposure that leadership can't audit and innovation that leadership can't scale.

The discipline of the Fast, Safe Yes is built on a different premise: the goal of guardrails is not to block adoption. It is to route adoption through a channel that is measurable, vetted, and aligned to organizational policy. When employees have a sanctioned path to surface the tools they're already using — and when that path is fast — they use it. When the alternative is a wall, they go around it.

What makes this work at scale is automation. Manual review of every AI tool request against NIST, FAR, the EU AI Act, and organizational policy is not sustainable at enterprise velocity. Automated vetting — where each tool request is evaluated against a defined policy framework and risk-tiered without human bottleneck — is how organizations get to yes faster while maintaining the controls they need.

WHAT THIS LOOKS LIKE IN PRACTICE

- A Safe Harbor program allows employees to surface tools they're already using privately, without fear of penalty — turning shadow usage into visible, manageable data
- Automated vetting evaluates each tool request against organizational policy, NIST, FAR, and the EU AI Act — without requiring a compliance team to review every submission manually
- Approved tools move into a governed registry with defined owners, KPI plans, and ongoing monitoring
- The governance foundation — an AI council with defined roles and decision rights — provides the institutional structure that makes the process credible and durable

73% reduction in AI-related security incidents for organizations with structured guardrails and automated policy enforcement. (Internal analysis)

66% of CEOs report that stakeholder trust concerns regarding AI, data, and security have arisen in the last 12 months. (PwC, 2026)

46% of organizations currently lack an official AI usage policy, while 95% use AI in some form. (General Assembly, 2026)

PART 4

The 5-Phase Loop: How Structured Adoption Works in Practice

The three disciplines above are not a one-time setup. They operate as a continuous cycle — each phase informing the next, and the loop improving with every pass. Here is what that cycle looks like in practice.

00

Foundation

Establish the AI Council, define roles and decision rights, and configure the organizational policy baseline — the framework that governs everything that follows. This is not a months-long project. An AI Council and foundational policy can be operational in fifteen minutes with the right tooling. The point is to start with structure, not retrofit it later.

01

Discovery

Set organizational AI goals, open the proposal marketplace for employee-driven ideas, and activate Safe Harbor to surface tools already in use. This phase transforms the typical "top-down mandate vs. bottom-up workaround" dynamic into a single, visible stream of AI activity that leadership can actually see and shape.

02

Vetting

Each proposed tool or initiative is automatically evaluated against organizational policy, NIST, FAR, and the EU AI Act. Risk-tiering determines what requires council review and what can be fast-tracked. The goal is not to slow adoption — it's to ensure that what moves forward is vetted, traceable, and aligned to the organization's risk tolerance.

03

Registry

Approved initiatives move into a governed registry — a real-time view of every sanctioned AI tool, its owner, its KPI plan, and its current performance. This is the single source of truth that leadership needs to answer the board question: what are we running, what is it costing, and what is it producing? Without this, AI program reporting is assembled from spreadsheets and memory. With it, it's a dashboard.

04

Readiness

For each tool that enters the registry, role-based enablement tracks are activated — training employees in the safe and effective use of that specific tool, for their specific role. Certification closes the loop: organizations don't just know which tools employees have access to. They know which employees are verified as capable of using them well. The cycle then feeds back to Discovery: new proposals are informed by what the workforce can actually support.

> **Each phase feeds the next. Outcomes are measured continuously, not after the fact. The loop is the product.**

PART 5

The Maturity Bridge: Where Are You Now?

Every organization that implements the framework above is moving along a maturity curve. The MITRE AI Maturity Model provides a standardized way to assess where an organization currently sits and what the path forward looks like.

The model runs from Level 1 (Initial — fragmented, experimental AI usage with no formal structure) through Level 5 (Optimized — AI embedded in the operating model with continuous improvement across people, process, data, and technology). Most enterprises today operate somewhere between Level 1 and Level 3: tools are deployed, some governance exists on paper, but the measurement and readiness disciplines are not yet operational.

The value of anchoring to this model is not the assessment itself — it's what the assessment reveals. Organizations that use a maturity framework can identify the specific gaps that are limiting their outcomes, sequence investments in the right order, and demonstrate progress to boards and stakeholders in terms that are standardized and defensible.

The 5-Phase Loop maps directly to the MITRE model. Foundation and Vetting address the governance and strategy dimensions. Readiness addresses workforce development and organizational structure. Registry and ongoing measurement address performance, data, and technology enablers. Moving through the loop is moving up the maturity curve — not as a separate initiative, but as a consequence of doing the work.

LEVEL	WHAT IT LOOKS LIKE	WHAT THE FRAMEWORK ADDRESSES
1 **Initial**	Ad-hoc experimentation. No policy, no measurement, no visibility.	Foundation (00): AI Council + Policy in place within days.
2 **Engaged**	Some tools deployed, policies drafted but inconsistently applied.	Discovery (01) + Vetting (02): Proposals captured, Safe Harbor live, automated vetting active.
3 **Defined**	Governance frameworks documented. Still manually enforced.	Registry (03): Approved tools in a single governed dashboard with KPI plans.
4 **Managed**	Outcomes tracked. Workforce training ad hoc.	Readiness (04): Role-based enablement tied to approved tools. Certification in place.
5 **Optimized**	AI embedded in operating model. Continuous improvement loop active.	Full 5-Phase Loop running. Program self-corrects. Leadership has real-time visibility.

CONCLUSION

From Drift to Direction

The organizations that will generate durable value from AI are not the ones that moved fastest. They are the ones that built the foundation to move well.

That foundation is not complex. It requires three disciplines — measuring before spending, enabling before deploying, and routing adoption through sanctioned channels rather than trying to block it. It requires a governance structure that knows how to say yes quickly and safely, not just how to say no. And it requires a measurement cycle that closes the loop between investment and outcome, so that what works gets scaled and what doesn't gets stopped.

The question every leader is facing right now is not whether to adopt AI. That question is already answered. The question is whether to lead that adoption or continue to be carried by it. Carried looks like more spend, more risk, and more activity that doesn't show up in the numbers. Leading looks like structure: defined outcomes, enabled people, and guardrails that accelerate rather than obstruct.

> **The gap between the two is not technical. It's organizational. And it's closable.**

Ready to move from speculation to structure?

We offer a 15-minute diagnostic to assess where your AI program currently sits on the maturity curve and identify the highest-leverage next step. No pitch. Just a structured conversation.

hello@clairet.com · clairet.com

APPENDIX

Source Index

All statistics cited in this paper are drawn from third-party research. Sources are listed below in alphabetical order by organization.

American Management Association (AMA). AI Becomes a Daily Workplace Tool with Employees Trying to Stay Ahead. AMA Global, December 2025.

Anthropic. How AI Assistance Impacts the Formation of Coding Skills. January 2026.

BCG Center for CFO Excellence. How to Get ROI from AI in the Finance Function. Boston Consulting Group, June 2025.

BearingPoint. AI Impact Enables Leaders, Employees and the Organization. February 2026.

DataCamp / YouGov. Companies Are Investing in AI, But Their Workforces Aren't Ready. Business Wire, February 2026.

Emburse / Talker Research. AI-Powered Finance: New Practices, ROI and Future Outlook. May 2025.

Ernst & Young. EY 2025 Work Reimagined Survey. November 2025.

General Assembly. The State of AI Adoption: By the Numbers. 2026.

LinearB. AI-Driven Code Generation is Not Doubling Productivity Without SDLC Transformation. February 2026.

Manufacturing Leadership Council / NAM. Shaping the AI-Powered Factory of the Future. May 2025.

McKinsey / QuantumBlack. The State of AI in 2025. November 2025.

METR. Measuring the Impact of Early-2025 AI on Experienced Open-Source Developer Productivity. July 2025.

MITRE Corporation. AI Maturity Model (AI MM). Version 1.0.

NBER. Firm Data on AI. Working Paper 34836. February 2026.

PwC. 29th Global CEO Survey: Leading through Uncertainty in the Age of AI. January 2026.

S&P Global Market Intelligence. Generative AI Experiences Rapid Adoption, but with Mixed Outcomes. May 2025.

The Conference Board. AI Risk Disclosures in S&P 500. 2025.

Writer. AI in the Enterprise: State of Adoption. 2025.