

CLEARFIELD ADVISORY

*Whitepaper*

# Before the Second Attempt

*A Structural Diagnosis of AI Rollout Failure — and What the Other 5% Understand*

---

**95% of enterprise AI pilots fail to deliver measurable financial return.**

42% of organizations abandoned most of their AI initiatives in 2025 alone — up from 17% the prior year.

The organizations now preparing a second attempt face a narrowing window and a question nobody has given them a precise answer to:

***What would need to be true about this organization for the second attempt to produce a different result than the first?***

**Sterling Welles Carroll**

Founder + Principal, Clearfield Advisory

*Restoring the Conditions for Organizational Joy*

clearfield-advisory.com

---

---

---

SECTION —

## Executive Summary

The failure of AI rollouts is not a technology problem. It is not a change management problem. It is not a training problem, a vendor selection problem, or a governance framework problem — though each of those has been offered, applied, and found insufficient at scale.

It is a structural problem. Specifically, it is the failure of five organizational conditions that must be sound before any implementation framework, workflow architecture, or adoption strategy can hold.

These five conditions — Authority, Orientation, Risk, Temporality, and Agency — form what Clearfield Advisory calls Human Decision Infrastructure. When this infrastructure is sound, AI multiplies organizational capacity. When it is fractured, AI multiplies organizational dysfunction. The technology performs precisely as designed in both cases.

The organizations in the successful 5% did not have better technology, more sophisticated governance frameworks, or more experienced change management partners. They had sounder infrastructure. They knew — explicitly or intuitively — who owned consequential AI decisions, what the organization was actually optimizing for, what the real cost of inaction was, how to align deployment timelines with adoption realities, and how to create the structural conditions for people to genuinely act from the tools they'd been given.

This whitepaper names that infrastructure precisely, maps the five failure expressions that appear when each element is fractured in an AI rollout context, and offers the diagnostic questions that reveal whether an organization's infrastructure is sound before the second deployment begins.

***The window for structural repair narrows between the first and second attempt. The organizations that address the infrastructure question before the second deployment have the full range of options available. Those that don't are working with a narrowing window at increasing cost.***

## SECTION 1

## The Pattern Nobody Is Naming

The numbers are no longer surprising. What is surprising is that they haven't produced a different diagnosis.

MIT's research, drawing on 150 executive interviews, 350 employee surveys, and analysis of 300 public AI deployments, found that 95% of enterprise generative AI pilots fail to deliver measurable financial return. S&P Global reports that 42% of organizations abandoned most of their AI initiatives in 2025 — nearly three times the abandonment rate of the prior year. RAND Corporation places overall AI project failure above 80%, twice the failure rate of non-AI technology projects. Bain & Company's analysis of more than 24,000 transformation initiatives found the 2024 failure rate at 88% — the highest recorded figure in three decades of tracking.

These are not marginal findings from niche research. They represent the preponderance of evidence from the most credible organizational research institutions in the world, converging on a single, consistent, and largely unexamined conclusion.

The failure is not random. It has a pattern. And the pattern points with remarkable consistency at a single structural reality that the current discourse has not yet named precisely.

***95% failure is not variance. It is signal. The question is what it's pointing at.***

Every major research source identifies the failure as organizational rather than technological. MIT specifically identifies 'flawed enterprise integration' and 'organizational learning gaps.' Forbes notes organizations fail because they treat AI as a project to oversee rather than a structural capability to embed. PwC's 2026 Global CEO Survey found that only 12% of CEOs report AI delivering both cost and revenue benefits — and attributes the gap to whether organizations have established 'strong foundations.'

Strong foundations. The phrase appears across the literature with remarkable consistency and almost no precision. What are the foundations? How are they assessed? How are they repaired when they're fractured? The research correctly identifies their absence as the cause of failure without naming what they are with sufficient specificity to be actionable.

This whitepaper names them.

## SECTION 2

## Why Current Solutions Don't Reach It

The organizational response to AI rollout failure has been consistent and, by the numbers, insufficient. When implementations stall or collapse, organizations reach for one of four familiar instruments. Each addresses something real. None addresses the structural layer where the actual failure lives.

**Workflow Governance** — The most technically sophisticated response to AI failure addresses how AI is integrated into existing processes at design-time and governed at runtime. The argument, well-made by practitioners in this space, is that AI should serve validated workflows rather than reason independently — reducing hallucination risk and producing consistent outcomes. This is correct as far as it goes.

*What it doesn't ask: who designed the workflow, and whose decision-making frame does it reflect? A perfectly designed workflow deployed into an organization where decision ownership is unclear, where the leadership team doesn't share a frame for what the AI is optimizing for, and where people have access to the tool but not genuine permission to act from it will produce the same inconsistency the workflow was designed to prevent. The workflow is the governance engine's most visible expression. The human decision infrastructure beneath it is the engine itself.*

**Change Management** — The behavioral response to AI failure addresses how people are prepared for transition, communicated with about change, and supported through adoption. The 70% organizational change failure rate that preceded the AI era should have been a signal that this instrument has limits. Applied to sound structural infrastructure, change management produces real results. Applied to fractured infrastructure, it absorbs investment and leaves the fracture intact.

*What it doesn't ask: why are the structural conditions that would allow change to hold absent — and what would it take to repair them before the intervention begins?*

**Vendor Selection and Technology Optimization** — The technology response to AI failure assumes the problem is in the model, the platform, or the implementation partner. Better technology, more carefully selected, more rigorously implemented. This response is expensive, time-consuming, and produces the second attempt's most common outcome: the same organizational failure pattern expressed through a more sophisticated technical substrate.

*What it doesn't ask: if the organization couldn't carry the first implementation structurally, what has changed in the organizational infrastructure that would allow it to carry a better one?*

---

Training and Upskilling — The human capital response to AI failure addresses whether people have the technical capability to use AI tools effectively. This is the most behaviorally sophisticated response and the most structurally incomplete. People can be trained indefinitely. If the structural conditions for Agency — clarity about what they're authorized to decide with AI, psychological permission to use it in consequential moments, cognitive bandwidth that isn't already consumed by the transition itself — are absent, the training produces capability without deployment.

*What it doesn't ask: what structural conditions would need to be present for trained people to actually act from their capability?*

Each of these responses addresses a floor above the structural failure. The failure is not in the technology, the change management practice, the vendor relationship, or the training program. The failure is in the five structural conditions that determine whether any of those instruments can hold.

---

### SECTION 3

## The Five Structural Conditions

Human Decision Infrastructure consists of five elements. They are not sequential — they are interdependent. A fracture in one element stresses the others. A fracture in multiple elements produces the compounding failure pattern that most AI rollouts eventually become.

What follows is each element named precisely, its specific failure expression in an AI rollout context, and the diagnostic question that reveals whether it is sound or fractured before the second deployment begins.

### **A Authority**

#### AI ROLLOUT FAILURE EXPRESSION

*Who owns the consequential AI decision — not who approved the deployment, not who chairs the governance committee, but who holds the decision when AI produces a wrong outcome and can stop it, change it, or own the result? In most failed AI rollouts, that question produces three different names, a governance chart that everyone knows doesn't reflect how decisions actually move, or silence. The silence is the fracture. AI doesn't create unclear Authority — it exposes it at a speed and scale the organization wasn't prepared for.*

## DIAGNOSTIC QUESTION

*"If I asked your five most senior leaders independently who holds the decision when your AI initiative gets something consequentially wrong — would they give the same name?"*

## **O** Orientation

## AI ROLLOUT FAILURE EXPRESSION

*What is this organization actually optimizing for with AI — and would your leadership team give the same answer independently? Finance sees cost reduction. Operations sees process automation. HR sees talent analytics. The CTO sees competitive infrastructure. The CEO sees transformation. All of those frames are legitimate. None of them are the same. The AI initiative launched into an organization whose leadership was optimizing for five different outcomes without knowing they differed. The initiative becomes a proxy war between frames rather than a coherent deployment. Orientation fracture is the most common invisible fault line in failed AI rollouts — and the least likely to appear in post-mortems.*

## DIAGNOSTIC QUESTION

*"If I asked each member of your leadership team to complete this sentence independently — 'what we are using AI to decide that we couldn't decide as well before is...' — how many different answers would I get?"*

## **R** Risk

## AI ROLLOUT FAILURE EXPRESSION

*The AI risk conversation in most organizations is entirely one-directional — focused on the risk of action. Hallucination. Data privacy. Regulatory exposure. Reputational cost. These are real risks and they deserve serious attention. What receives almost no serious attention is the risk of inaction — the compounding cost of the structural gap while the organization debates whether to address it. For organizations preparing a second AI attempt, this cost has already been accumulating since the first initiative stalled. Every week of unmeasured inaction has a price. The inability to name it is itself diagnostic.*

## DIAGNOSTIC QUESTION

*"What has the gap between your first AI initiative's promise and its actual outcome cost this organization — in resource misallocation, in lost competitive position, in the organizational patience that won't be available for a third attempt if the second one fails?"*

---

## T Temporality

### AI ROLLOUT FAILURE EXPRESSION

*Every AI initiative has a deployment timeline. Milestones, phases, go-live dates — all tracked with precision. What almost no AI initiative has is an adoption timeline. Because deployment and adoption run on different clocks. Deployment runs at technology speed — fast, sequential, milestone-driven. Adoption runs at human infrastructure speed — slower, non-linear, entirely dependent on whether the structural conditions for genuine use were built before the tool arrived. When those two clocks are unsynchronized — and they almost always are — the organization declares the rollout complete while the actual work of making it real hasn't begun. The milestone was the deployment. The outcome was supposed to be the decisions.*

### DIAGNOSTIC QUESTION

*"On the day your first AI initiative went live — what had changed in the five structural conditions that would allow your people to actually decide from it? Not the technology. The structural conditions."*

---

## Ag Agency

### AI ROLLOUT FAILURE EXPRESSION

*The average AI initiative that reaches production achieves 12% adoption. Not because the tool doesn't work. Not because people are resistant to change. Not because the training was inadequate. Because the people who were supposed to use it didn't have the structural conditions to actually act from it. Clarity about what they were authorized to decide with it. Psychological permission to use it in consequential moments — not just low-stakes ones. Cognitive bandwidth that wasn't already consumed by the transition the AI rollout was layered on top of. Agency is the practical capacity to act — authority, clarity, permission, and bandwidth combined. When Agency is absent, the tool sits there. Expensive. Functional. Unused.*

### DIAGNOSTIC QUESTION

*"Who in your organization has formal access to the AI tool but has stopped using it — or never started — and do you know the structural reason why?"*

---

These five elements are not a checklist. They are an interdependent architecture. Authority fracture produces Orientation confusion — when nobody clearly owns the AI decision, different parts of the organization develop different frames for what it's for.

---

---

Orientation fracture produces Agency collapse — people can't act from a tool whose purpose they don't share with their leadership team. Temporality misalignment produces Risk accumulation — the deployment clock moves while the adoption conditions remain unbuilt, and the cost compounds silently.

The organizations in the successful 5% are not distinguished by having solved each element independently. They are distinguished by having maintained the structural coherence between all five — before the deployment began.

---

#### SECTION 4

## The Overton Window Closes

There is a concept in political science called the Overton Window — the range of ideas a society will currently accept as reasonable. Outside the window: too radical, too unfamiliar, not yet thinkable. Inside the window: debatable, acceptable, actionable.

Structural fault lines in organizations have their own Overton Window. And it closes faster than most organizations realize.

When a structural fracture first appears — when Authority becomes unclear, when Orientation begins to diverge, when Agency starts to withdraw — the full range of repair options is available. The fracture is recent. The workarounds haven't calcified. The people closest to the problem haven't yet reorganized their behavior around it. The window is open.

***Then the organization adapts. And the window begins to close.***

Workarounds normalize. The person who knew how to navigate the unclear Authority structure becomes indispensable — and protective of the knowledge. The teams optimizing for different AI outcomes develop siloed workflows that reflect their divergent Orientation. The people who stopped using the AI tool find alternative processes that don't require it — and quietly advocate for those processes to continue.

The political cost of naming what's actually wrong increases as more people invest in the structure that replaced the original one. The post-mortem that would have been difficult six months ago becomes organizationally dangerous twelve months later.

By the time most organizations correctly identify the structural fracture — which is

---

---

typically twelve to eighteen months after it first appeared — the window has narrowed significantly. Not because the repair is impossible. Because the organization has reorganized itself around the fracture in ways that make the repair more complex, more expensive, and more politically costly than it would have been at the point of origin.

For organizations preparing a second AI attempt, this is the critical question the Overton Window poses:

***How much has the window narrowed since the first attempt failed?***

The structural conditions that produced the first failure didn't pause while the organization decided whether to try again. They continued to develop. The Authority fracture that produced unclear decision ownership in the first rollout has had another twelve to eighteen months to calcify into informal workarounds. The Orientation divergence that produced competing frames for what AI was for has had time to harden into competing departmental strategies. The Agency collapse that produced 12% adoption has had time to become the organizational assumption that 'people here just don't use AI.'

The second attempt that proceeds without addressing the infrastructure is not starting fresh. It is starting from a more fractured position than the first attempt began — with less organizational patience, less board goodwill, and a narrower window for repair.

This is not an argument against second attempts. It is an argument for addressing the structural conditions before the second deployment begins rather than after it stalls.

---

SECTION 5

## **Before the Second Attempt: Five Questions**

The Decision Integrity Diagnostic begins before an engagement. It begins with the questions a leadership team cannot answer cleanly — or cannot agree on. The five below are drawn directly from the diagnostic instrument Clearfield Advisory uses in structural assessments. They are not rhetorical. They are diagnostic.

A leadership team that can answer all five cleanly and consistently has sound Human Decision Infrastructure. The second deployment will hold.

---

A leadership team that cannot — or that produces divergent answers across its members — has located the structural work that needs to happen before the next initiative begins.

**Authority** — When the last consequential AI decision stalled or produced an unexpected outcome, who actually held it — and would everyone in the leadership team give the same answer?

**Orientation** — If you asked each member of your leadership team independently what this organization is using AI to decide that it couldn't decide as well before — would the answers be the same?

**Risk** — Can you name the specific AI-related decision most likely to degrade in execution over the next ninety days — and who owns it?

**Temporality** — Which decisions made during the first AI initiative are still being executed — even though the conditions that justified them no longer exist?

**Agency** — Who in your organization has formal access to AI tools but has stopped using them — or never started — and do you know the structural reason why?

These questions do not require an engagement to be useful. They require honest conversation among the people who would need to answer them together. If that conversation has not happened — or has stalled — the Fault Map will show why.

The organizations that do this work before the second deployment don't just improve their odds of technical success. They change the structural conditions that determine whether any technical success is possible. Their people stop being given access — which is compliance — and start having genuine Agency — which is the practical capacity to make consequential decisions from the tools they've been given.

We call the outcome of that structural repair organizational joy. Not as aspiration. As structural outcome. The condition in which capable people, operating inside sound infrastructure, can finally act from their full capacity rather than adapting to a structure that has been working against them.

SECTION —

## A Final Note

The 5% of organizations whose AI initiatives deliver measurable financial return are not distinguished by superior technology, more sophisticated governance, or more experienced implementation partners. They are distinguished by something more fundamental and more replicable: they understood the structural conditions their organizations needed to carry the initiative before the initiative began.

That understanding is not proprietary knowledge. It is not a competitive advantage that accrues only to large enterprises with dedicated AI governance teams and unlimited implementation budgets. It is a structural assessment — precise, contained, and actionable — that any organization preparing a second AI attempt can commission before the second deployment begins.

The window is open. The question is whether the organization is willing to look at what the first attempt was actually trying to tell it.

***The failure rate isn't the problem. It's the diagnosis the failure rate has been trying to surface for three decades.***

Clearfield Advisory works with organizations in transition — specifically those preparing a second AI attempt — to map the structural fault lines that behavioral interventions cannot reach, and to create the conditions for their leadership teams to repair them.

The repair belongs to the organization. We provide the map and the working session instrument that allows the leadership team to design the repair themselves, in a single session, without ongoing practitioner dependency.

---

### **Sterling Welles Carroll**

Founder + Principal, Clearfield Advisory

[clearfield-advisory.com](http://clearfield-advisory.com)

*Restoring the Conditions for Organizational Joy*

---

SOURCES REFERENCED

*MIT Project NANDA — Analysis of 150 executive interviews, 350 employee surveys, and 300 public AI deployments. Fortune, August 2025.*

*S&P Global — AI Project Abandonment Survey, 2025. CIO Dive, 2025.*

*RAND Corporation — AI Project Failure Rate Analysis, 2024–2025.*

*Bain & Company — Analysis of 24,000+ transformation initiatives, 2024.*

*PwC — 2026 Global CEO Survey, January 2026. Survey of 4,454 CEOs across 109 countries.*

*Gallup — Change Management Index for Q12®, 2016.*

*McKinsey Global Surveys — Organizational Transformation Failure Rate, 2008–2024.*

**CLEARFIELD ADVISORY**

© 2026 Clearfield Advisory. All rights reserved.