

The AI Savings Never Showed Up. The Security Debt Did.

A note from someone who's been doing this for 29 years, on why experience matters and guidelines matter.

Daniel Rodriguez • Founder & CEO, Infragil • September 2026

Everyone is using AI. Almost nobody can prove it's making them money.

That's not me being a grumpy security guy. McKinsey's latest State of AI survey says 88% of organizations now use AI in at least one business function. Sounds great. Keep reading: only 23% are scaling agentic AI anywhere in the enterprise, nearly two thirds haven't begun scaling AI across the enterprise at all, and only about 6% qualify as high performers seeing real bottom-line impact.

PwC asked 4,454 CEOs an even simpler question: is AI actually making you money? Only 12% said it delivered both revenue gains and cost savings. 56% reported no significant financial benefit at all. And 22% said their costs actually went up.

Then there's Gartner, predicting that over 40% of agentic AI projects will be canceled by the end of 2027. Not because the models don't work. Because of escalating costs, unclear business value, and inadequate risk controls.

Read that last phrase again. Inadequate risk controls. We'll come back to it.

And before anyone files this under "old guy dunking on AI," stop. I run an AI-native company. I use these tools every single day, and they're the biggest force multiplier I've seen in three decades of doing this work. AI isn't failing. The business implementation approach is failing. Cut the experienced people first, hand the tools to whoever's left, skip the training, skip the governance, wait for savings. That's the play most companies ran, and it's why 56% have nothing to show for it. The 12% who are winning aren't using better models. They're implementing with better people, better preparation, and better guardrails. The rest of this article is the evidence, and then the playbook.

The layoff boomerang

Companies didn't wait for the ROI to show up before cutting people. Roughly 55,000 U.S. job cuts were blamed directly on AI in 2025, and the pace picked up into 2026. The pitch was simple math: AI replaces headcount, headcount is cost, so AI equals savings.

Except the savings never materialized. And now the rehiring has started.

Forrester found 55% of employers already regret their AI-driven layoffs. CareerMinds surveyed HR leaders who ran AI-related cuts: two thirds are already rehiring, and over half started within six months

of the original layoffs. Six months. Gartner projects that by 2027, half the companies that blamed AI for headcount reductions will be rehiring for similar functions.

Want the math on that round trip? Nearly 31% of organizations spent more on rehiring than the layoffs saved. Another 42% broke even. Orgvue found companies spend about \$1.27 for every \$1 "saved" through workforce cuts once you count severance, lost productivity, and replacement costs.

Klarna said its chatbot did the work of 700 agents, watched satisfaction slide, and reversed course. IBM automated most routine HR requests, then announced plans to triple entry-level hiring, because as their CHRO put it, if you don't invest in the pipeline, the well dries up.

So we fired the people, kept the technology that couldn't replace them, and now we're paying a premium to buy the judgment back. That's not transformation. That's a very expensive lesson in what judgment was worth.

I've seen this movie. It was called "the cloud."

If this feels familiar, it should. I lived through the first run.

Fifteen years ago the pitch was almost word for word the same: move to the cloud, cut infrastructure costs, cut headcount. Companies rushed in without the architecture expertise, without cost governance, without security fundamentals. The bill is still being paid.

Five years ago, organizations estimated 30% of their cloud spend was pure waste, with budgets blown by an average of 24%. Today, after fifteen years of maturity and an entire FinOps industry invented specifically to fix this, wasted spend still sits at 29%. It actually went up last year. In Flexera's 2026 report, 85% of organizations named managing cloud costs their top challenge. Ahead of security. IDC found 59% of organizations blew their cloud budgets in 2024.

The refactoring never stopped either. The Barclays CIO Survey found 83% of CIOs planning to repatriate at least one workload, up from 43% in late 2020. A Citrix survey found 93% of IT leaders had pulled some workloads back out of the cloud over three years. Think about that. Companies are unmigrating systems they spent years migrating, because the "lift and shift" shortcuts taken in the rush never fit the platform.

And some companies didn't survive the lesson. In 2014, Code Spaces, a code-hosting provider that promised customers "rock solid, secure" infrastructure, went from a viable business to permanently shut down in about 12 hours. An attacker got into their AWS control panel. When the company tried to take back control, the attacker deleted their data, their machine configurations, and their backups. Everything, including the backups, lived under one poorly protected admin console. No MFA. No privileged credential management. People called it the first death in the cloud and said it would be a case study forever. It wasn't exotic. It was the basics, skipped in the excitement of a new platform. The exact basics people like me had been preaching for decades.

The cloud wasn't the problem, and neither is AI. The cloud is a phenomenal technology when it's architected by people who understand it, governed by people who measure it, and secured by people who've seen how it fails. The companies that won in the cloud didn't avoid it. They brought their experienced people with them.

New platform, breathless savings promises, expertise treated as friction, then a decade of overspend, refactoring, and breaches. AI is running the same play. Just faster. And this time the platform writes its own insecure code.

The "prompt engineer" problem

Somewhere along the way a strange idea took hold: that a weekend of prompting makes someone an engineer, an analyst, a security architect. That "AI expertise" replaces domain expertise instead of multiplying it.

So companies started trading educated, experienced professionals for people who are fluent at asking a model for things but have no background to judge what comes back.

Here's the problem with that trade. AI is extraordinary at producing things that look finished. Clean UI. Working features. Confident code. What it doesn't reliably produce, especially for someone who can't check its work, is things that are safe.

The numbers on AI-generated code are rough:

- Roughly 62% of AI-built applications ship with critical security vulnerabilities. They work perfectly. They're just not secure.
- Apiiro looked at Fortune 50 enterprises and found AI-generated code had 322% more privilege escalation paths and 153% more design flaws than human-written code. Those are architectural failures. Scanners don't catch them and novices never see them.
- Retool surveyed 307 CTOs, CIOs, and CISOs: 22% of organizations already had a production incident caused by an AI-generated internal tool. Another 51% can't say for certain whether they have. Read that again.
- Gartner is warning that "prompt-to-app" development by citizen developers will increase software defects by 2,500% by 2028. Their words: a software quality and reliability crisis.

When the finance analyst ships an app

This scenario is playing out in thousands of companies right now, and I want to be clear that the people in it aren't the villains.

A sharp finance analyst, genuinely smart, genuinely well-intentioned, uses an AI tool to build a slick internal web app for tracking vendor payments. It works. Leadership loves it. It gets connected to real data and real users.

Nobody asked about input validation. Nobody reviewed the authentication logic the model generated. Nobody noticed the API key sitting in client-side code, or the database query that's one crafted input away from an injection attack. The analyst didn't fail. The analyst was never equipped to spot injection flaws, over-broad permissions, or unsafe data handling. That was never their job.

Or take the assistant who gets told "build something with AI, no guidance needed." The model helpfully wires up remote access so it's easy to update from home. Default credentials. Exposed port. What just got built, functionally, is the foothold ransomware operators pay for. Nobody was malicious. There was just nobody with security experience in the loop, and no guideline that required one.

That's shadow AI in real life. Gartner found over 57% of employees use personal GenAI accounts for work, and a third admit putting sensitive data into unapproved tools. Your attack surface isn't growing in the IT department. It's growing everywhere except IT.

McKinsey found 51% of firms have already had AI-related incidents. The ones handling it well share a profile: human-in-the-loop rules, centralized oversight, executive accountability. Governance. Experience. Guardrails.

Experience isn't gatekeeping. It's the enabler.

There's a reason architect and engineers stamp drawings. A reason physicians carry licenses. A reason security people spend decades learning how systems fail before anyone trusts them with the systems that can't. Experience isn't a certificate on a wall. It's pattern recognition built from everything you've watched go wrong. AI didn't eliminate the need for it. AI made it the highest-leverage investment in your whole program, because one experienced person can now enable, review, and safeguard the AI-amplified output of an entire team.

The model that works is simple: experienced people enabling AI. Not AI replacing experienced people.

PwC's data backs this up. The 12% who are winning built foundations first: data quality, integration, formal responsible AI and risk management, before they scaled. The blockers haven't changed in thirty years. Data. Integration. Security. Governance. Adoption. Those are people-and-process problems, and they get solved by people who've solved them before.

The playbook

If your AI initiative is stalled, stealthy, or scary, here's where I'd start:

- 1 Put experienced people in charge of enablement.** Your veterans aren't the obstacle to AI adoption; they're the delivery mechanism. Give senior engineers, architects, and security leaders the mandate and the time to build the paved road: approved tools, reference patterns, secure defaults. Adoption led by the people who know where the bodies are buried survives contact with production.
- 2 Train the people you have.** The companies rehiring at \$1.27 on the dollar are paying for "tech to replace humans" budgets that never included upskilling. Flip it. Budget training alongside tooling. A workforce that understands your business and knows how to use AI safely is exactly the asset your competitors are trying to rebuild right now.
- 3 Hire qualified people and stop being "prompting" as a qualification.** Prompt fluency takes a week. Judgment takes a career. Hire for the underlying discipline (engineering, security, finance, operations) and let AI multiply it.
- 4 Never hand AI to an entry-level person without a mentor.** This one will cost you more than any tool ever saves. An entry-level person with AI and no mentor produces senior-looking output with junior-level judgment, and there's nobody in the loop who can tell the difference. That's exactly how the vulnerable app ships and the backdoor get built. Pair every junior AI user with someone experienced enough to review what comes out. Mentorship isn't a perk anymore. It's security control. IBM figured this out the hard way; it's why they're tripling entry-level hiring and rebuilding the pipeline instead of cutting it.
- 5 Find what's already been built, then hold AI code to the same bar as human code.** You can't govern what you can't see, and shadow AI apps touching real data are already in your environment. Inventory them. Then apply the same review, testing, least privilege, and secrets management you'd demand of anything else. No exceptions for "it's just an internal tool."
- 6 Fund governance like it's the product.** Gartner already told you inadequate risk controls will kill 40% of these projects. Governance isn't the brake. Done right, it's what lets you go fast without going through the windshield.

AI didn't make experience obsolete. It made experienced people the difference between the 12% and everyone else.

There are reasons experience matters. Reasons guidelines matter. We've paid this tuition before. We don't have to pay it again.

At Infragil, we help companies use AI responsibly and effectively. We put the experienced leadership, training, governance, and security foundations in place that separate the organizations getting real

value from the ones headed for the cancellation statistics. If your team is somewhere between "everyone's using AI" and "nobody can prove it's safe or paying off," that's the gap we close. Let's talk.

Daniel Rodriguez is the Founder & CEO of Infragil, an AI advisory and cybersecurity practice, and creator of Skopos, an AI-native third-party risk management platform. He has spent 29 years in cybersecurity, including CISO and VP roles.

Sources

- McKinsey, The State of AI in 2025 (adoption, agentic scaling, high performers, AI incidents)
- PwC, 29th Global CEO Survey (Jan 2026): AI revenue/cost outcomes across 4,454 CEOs
- Gartner (Jun 2025): 40%+ agentic AI project cancellations predicted by end of 2027
- Gartner: 2,500% defect increase from citizen-developer prompt-to-app by 2028; shadow AI usage
- Challenger, Gray & Christmas: 2025 AI-attributed job cuts
- Forrester, Predictions 2026: 55% employer regret; layoff reversals
- CareerMinds (Feb 2026): rehiring rates and costs after AI layoffs
- Orgvue: \$1.27 cost per \$1 saved in workforce reductions
- Apiiro: privilege escalation and design flaws in AI-generated code
- Retool, The State of AI Governance in 2026: production incidents from AI-built tools
- OX Security: vulnerability rates in AI-built applications
- Flexera, State of the Cloud (2021 to 2026): wasted cloud spend, budget overruns, cost management challenges
- IDC: 2024 cloud budget overspend; repatriation rates
- Barclays CIO Survey: workload repatriation intent (83% in 2024 vs. 43% in 2020)
- Citrix (2024): IT leaders moving workloads back from cloud
- Contemporary reporting on the 2014 Code Spaces breach and shutdown (Threatpost, Dark Reading, The Register, Network World)