JOYRIDE VENTURES

REVENUE STRATEGY

The Revenue AI Gap

Why the teams with the most to gain from AI are capturing the least of it



Chris Jensen
Principal, Joyride Ventures · 5 minute read

IN BRIEF

The problem

Sales and marketing stand to gain more from generative AI than any other function. Yet inside the same companies, revenue teams adopt AI at a fraction of the rate engineering does.

Why it happens

Three structural failures: leaders lack hands-on fluency, every AI session starts from zero context, and distrust of AI output is rational in a function where one error can cost a deal.

The fix

A Crawl, Walk, Run sequence: build a shared knowledge base, standardize high-leverage workflows, then connect agentic systems to the systems of record. The CRO becomes the architect of that system.

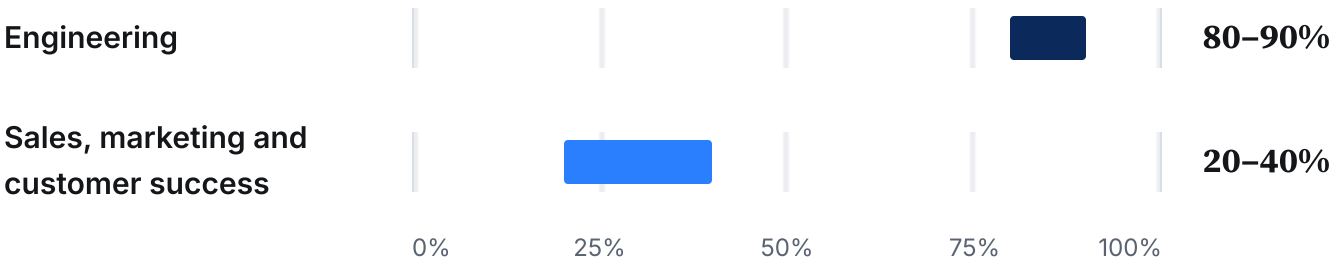
McKinsey estimates that generative AI could unlock up to $1.2 trillion in productivity for sales and marketing, more than any other function in the enterprise. And yet, by virtually every measure that matters, revenue organizations are among the slowest to claim it.

That is the central paradox in enterprise AI. According to Worklytics and Larridin's 2026 AI benchmarks, engineering teams at mature companies are running at 80 to 90 percent meaningful AI adoption. Sales, marketing, and customer success at the same companies, with the same leadership and the same budgets, are operating at 20 to 40 percent. Same company. Four times the adoption rate.

EXHIBIT 1

The same company, two different adoption curves

Share of teams with meaningful AI adoption at mature companies



Source: Worklytics and Larridin, 2026 AI benchmarks

The default explanation is cultural resistance. I do not think that fully explains the issue. After sitting with more than a dozen CROs and CEOs over the past year, I am convinced the issue is not unwillingness. It is structural. Three organizational failures, operating simultaneously, make meaningful adoption a persistent uphill battle even for willing teams.

The fluency problem runs from the top down

The first wave of AI tools was built for content creation and information retrieval, which is exactly the work senior leaders have largely delegated. But the deeper issue is that engineers have spent years building with these tools, breaking them, integrating them into real personal and professional workflows, and developing genuine intuition for where they help and where they fail. Most GTM professionals have not. They were late to the party, and their experimentation has been limited to basic tasks like email generation. That is not a recipe for fluency. It is a recipe for avoidance.

Context debt is the hidden tax on every AI interaction

Every morning, a sales rep opens a chat window and re-explains the company's ICP, competitive positioning, and quarterly priorities just to get a usable first draft. When the session ends, the context evaporates. The next rep starts from zero. So does the same rep the following morning. The AI is like a perpetual intern who shows up every Monday with no memory of last week.

This creates context debt, and it compounds. It explains why so many organizations report genuine time savings from AI while seeing no movement in win rates or pipeline quality. The activity metrics improve. The outcome metrics do not.

> *The activity metrics improve. The outcome metrics do not.*

Distrust is rational, not irrational

Varicent's 2025 research found that 44 percent of revenue leaders cite their own skepticism as a greater barrier to AI value than any technical issue. That is often viewed as a mindset problem. I read it differently. A hallucinated detail in a client brief can derail a six-figure deal. A miscalibrated forecast triggers the wrong resource decisions upstream. The sensitivity to AI error in revenue is not conservatism. It is appropriate risk calibration.

The problem deepens when AI becomes a substitute for thinking rather than an accelerant for it. Reps who consistently outsource their analytical work stop developing the critical thinking that made them effective in the first place, including pattern recognition, objection handling, and deal signal instincts. When AI does that work, the rep gets the output without building the skill. That is not a technology failure. It is a discipline failure, and it compounds quietly until the damage is already done.

A framework for breaking the loop

My clients who have made the most meaningful progress closing this gap follow a three-phase approach: Crawl, Walk, Run.

EXHIBIT 2

Crawl, Walk, Run

Each phase builds on the infrastructure of the one before it

01	02	03
Crawl	**Walk**	**Run**
ELIMINATE CONTEXT DEBT	**STANDARDIZE THE WORKFLOWS**	**CONNECT AGENTIC SYSTEMS**
A shared knowledge base with the ICP, competitive positioning and objection library that AI draws from automatically.	AI-assisted routines for high-leverage, repeatable tasks: pre-call research, account scoring, post-call follow-up, churn flags.	Multi-step workflows tied to systems of record that update deal stage and forecast category and surface deal health signals.

Crawl eliminates context debt. The output is a shared institutional knowledge base, a living repository containing the organization's ICP, competitive positioning, and objection library that AI draws from automatically rather than requiring reps to reconstruct it from scratch each session. After building this system, one of my clients reduced manual meeting prep time by roughly 60 percent. The bigger win was that bottom-half AEs started showing up to calls as well-prepared as top-half AEs. AI did not raise the ceiling. It raised the floor.

> *AI did not raise the ceiling. It raised the floor.*

Walk standardizes AI-assisted workflows for high-leverage repeatable tasks: pre-call research, account scoring, post-call follow-up, churn flags. Not glamorous work. But in organizations where the top 17 percent of reps generate 81 percent of revenue, as found by Ebsta and Pavilion, raising the floor has a massive impact on productivity.

Run is where agentic applications enter: multi-step workflows that connect the shared brain and skill layer to your systems of record. A well-designed Run-phase system might ingest call transcripts, update deal stage and forecast category automatically, and surface deal health signals to the rep before they think to look. Teams may employ a live client ‘stock ticker’ that continuously surfaces account health to key stakeholders. Published case studies, including from LangChain, at this level show 250 percent increases in conversion rates. But they are built

on top of functioning Crawl and Walk infrastructure. A word of caution: connecting the systems is the easy part. The real work is in the data architecture underneath, and most organizations underestimate it significantly.

The CRO as Architect

The traditional CRO archetype is a motivator, a talent developer, a pipeline quarterback. That job description is not obsolete. But the challenge described here is not a people-management problem. It is an information architecture problem.

The CRO who succeeds in this environment is the architect of the shared brain. Their most consequential decisions are no longer about headcount, territory design, or commission structure. Those are table stakes. Instead, they are about what context their team operates from, what workflows they run, and whether the system they are building is measurably smarter every quarter than it was the quarter before.

That shift, from managing people to architecting the environment in which people and AI work together, is what separates the next generation of revenue leaders from the last one.



Chris Jensen served as Chief Revenue Officer of Ibotta through its NYSE listing. He advises CEOs and CROs of high-growth technology companies on revenue strategy and AI adoption through Joyride Ventures.

SOURCES REFERENCED

- McKinsey Global Institute, research on the economic potential of generative AI
- Worklytics and Larridin, 2026 AI adoption benchmarks
- Varicent, 2025 research on revenue leaders and AI
- Ebsta and Pavilion, B2B sales benchmark research
- LangChain, published customer case studies