Apollo

Apollo's 2026 AI in Sales & Go-to-Market Survey

About the Report

Apollo's 2026 AI in Sales & Go-to-Market Survey was conducted in June 2026, with 328 GTM professionals (customers and non-customers) responding.
This is the first edition of a research series designed to benchmark AI adoption across go-to-market teams.

August 2026

Ask ten GTM leaders if their team has adopted AI, and ten will say yes.

Ask those same ten what "agentic AI" means, and you could get ten different answers.

Go-to-market teams are universally adopting AI, and fast. What's still catching up is everything beyond the surface level: how mature that adoption actually is, what's blocking teams from going further, and how those blockers shapeshift depending on whether you're a team of five or five-thousand.

But the majority of sales and marketing leaders don't think AI is coming for their headcount. In fact, **only 6% think it could replace someone on their team today**. Instead, they're betting on AI to do the parts of the job that used to consume reps' day, such as finding prospects, researching accounts, and personalizing outreach, so reps spend more time actually selling.



These are the five key themes we found in AI adoption across GTM teams.

AI adoption across GTM teams is universal — understanding and maturity aren't.

97%

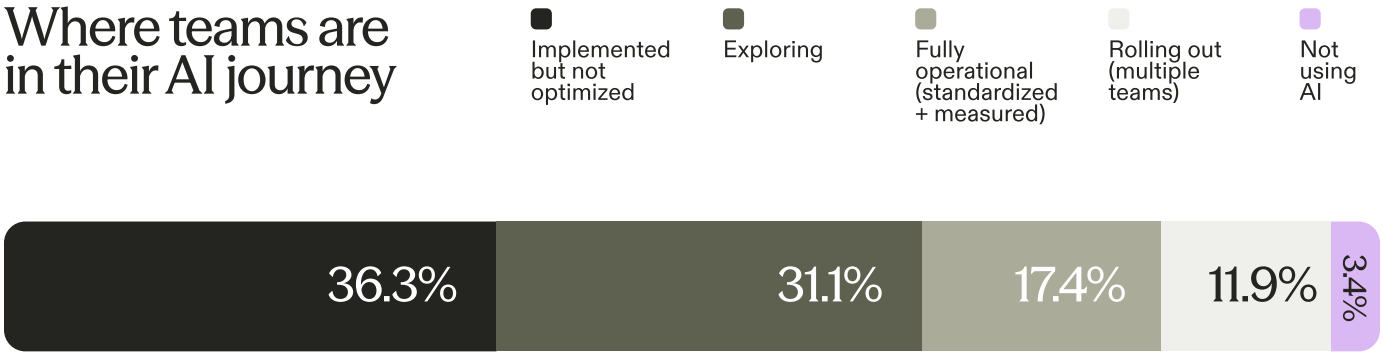
of GTM teams have adopted AI

97% of GTM teams have adopted AI.

That number stopped being interesting the moment adoption stopped being a question.

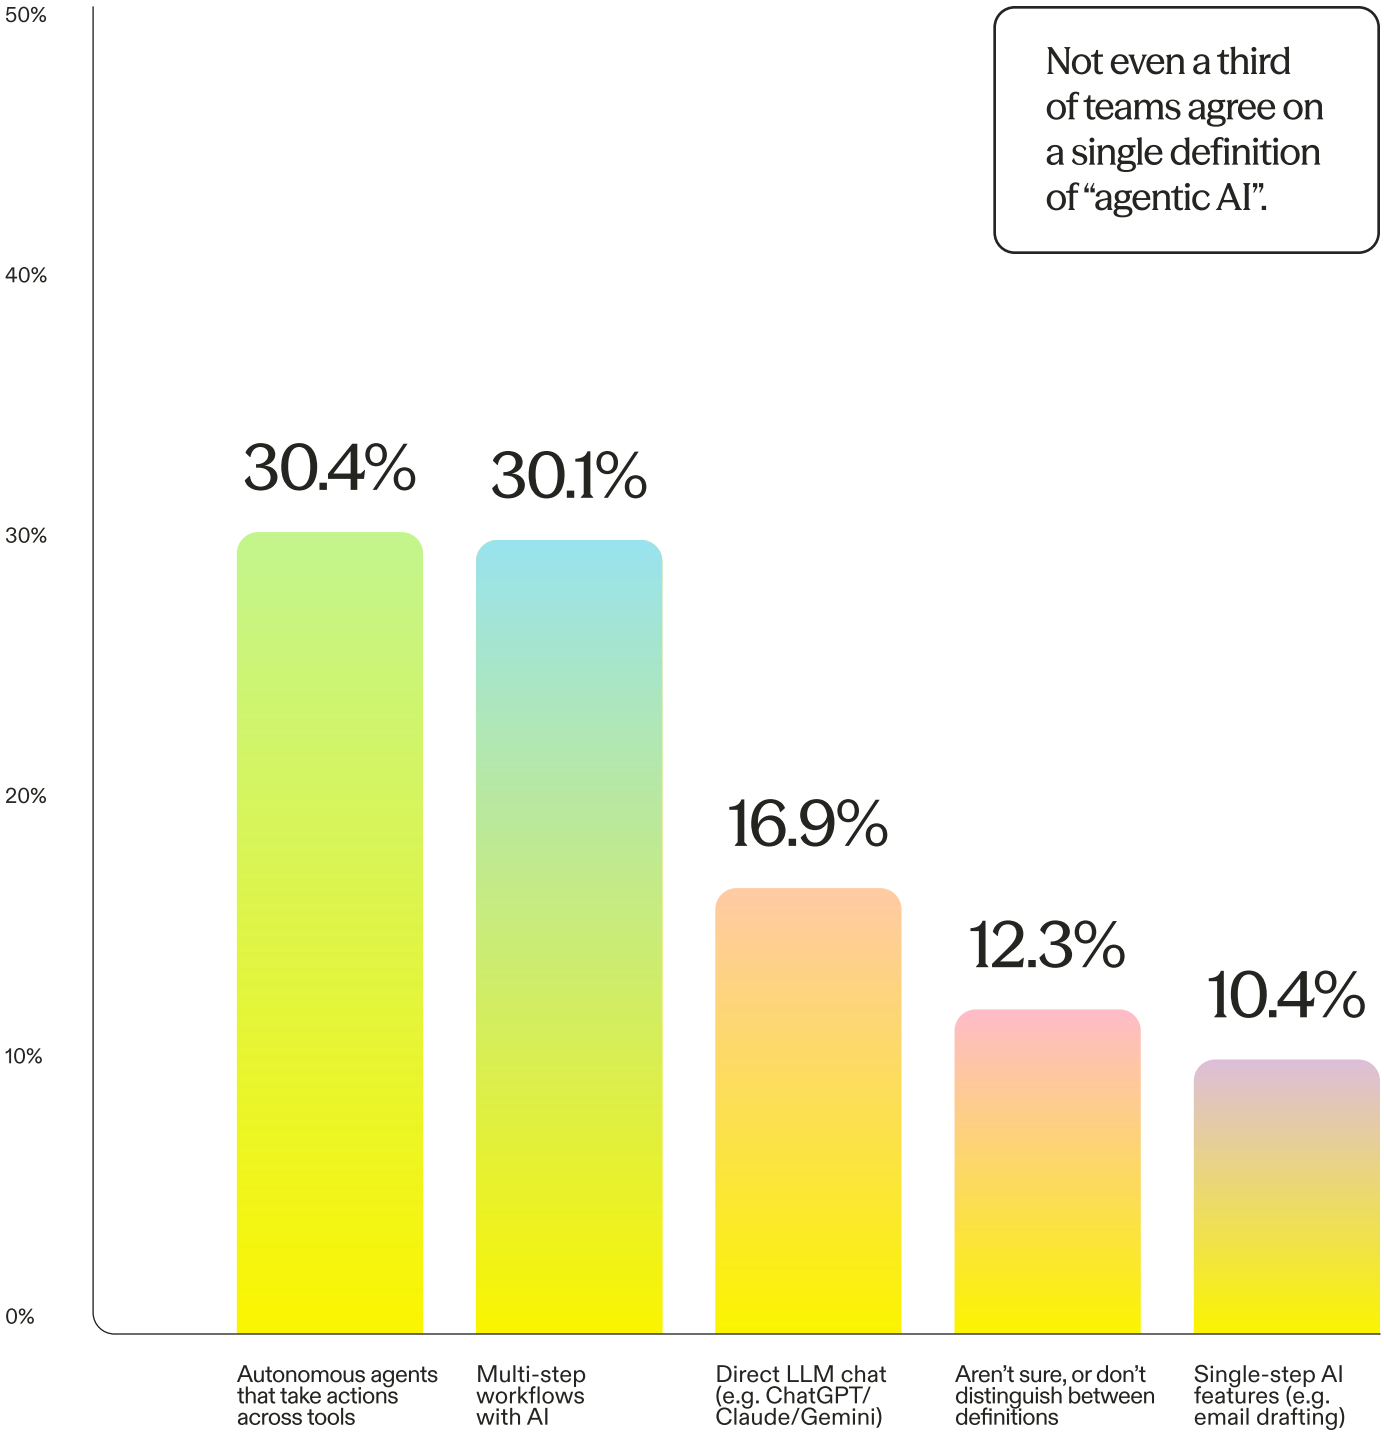
The real story is what's happening after a team implements AI. Our research shows that while most teams are using AI, they're far from optimized. They've moved past just exploring, but haven't reached anything a leader would confidently call mature.

Where teams are in their AI journey



That's a wide spread for one technology category — wide enough to ask whether "agentic AI" even means the same thing to everyone claiming it. No surprise here: it doesn't. **Not even a third of teams agree on a single definition.**

What “agentic AI” means to GTM teams



Not even a third of teams agree on a single definition of “agentic AI”.

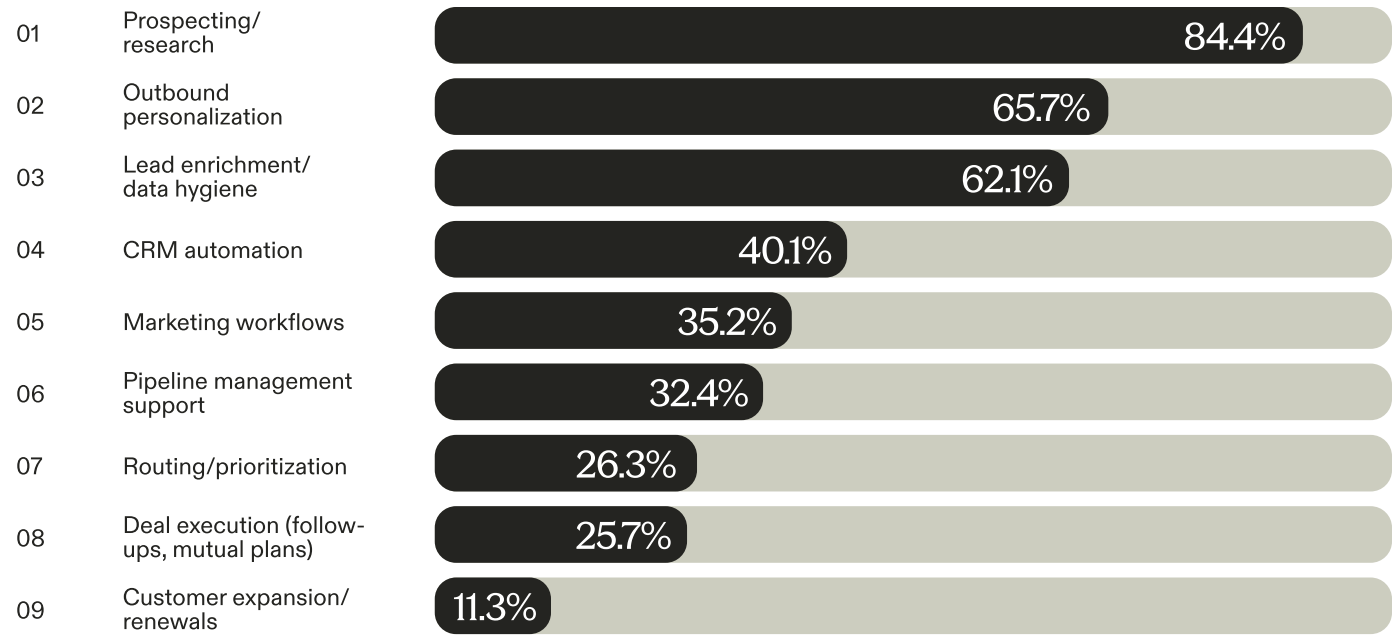
Prospecting leads AI adoption by a landslide.

Ask any GTM team where AI shows up first, and the answer is nearly universal: prospecting and research.

No matter the company size, from the smallest teams to the largest, an overwhelming majority call it their #1 use case. Teams reach for AI to find and research prospects long before they trust it anywhere else in the funnel.

AI use doesn't stop at prospecting for every team, but for most, that's where the momentum starts to fade. It holds up through outbound personalization and lead enrichment/data hygiene, then drops off fast. By deal execution and renewals, AI use drops significantly.

GTM AI use cases, ranked by adoption

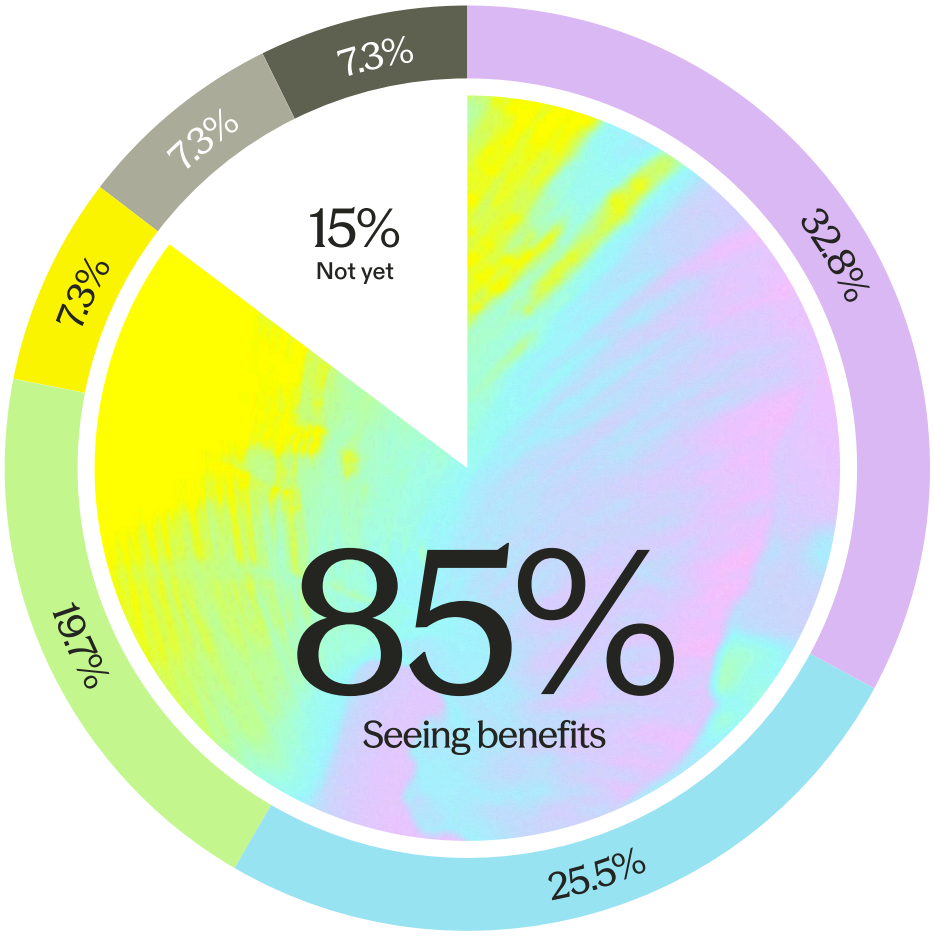


AI proves itself fast, and teams are confident it'll keep paying off.

Nearly a third of GTM teams using AI saw results in under 30 days and more than half saw them within two months.

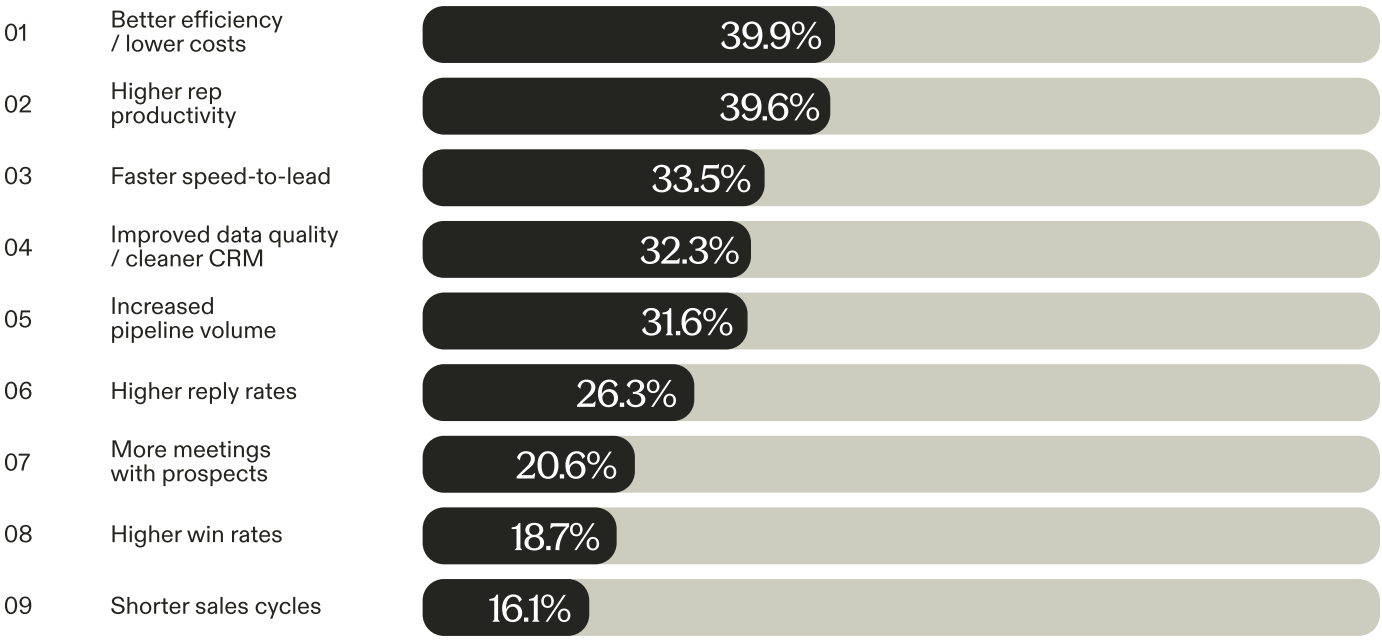
How long it took to see measurable benefit

- >30 days..... **32.8%**
- 1 – 2 months..... **25.5%**
- 3 – 6 months..... **19.7%**
- 6+ months **7.3%**
- Not seeing benefits yet **7.3%**
- Haven't implemented AI workflows..... **7.3%**



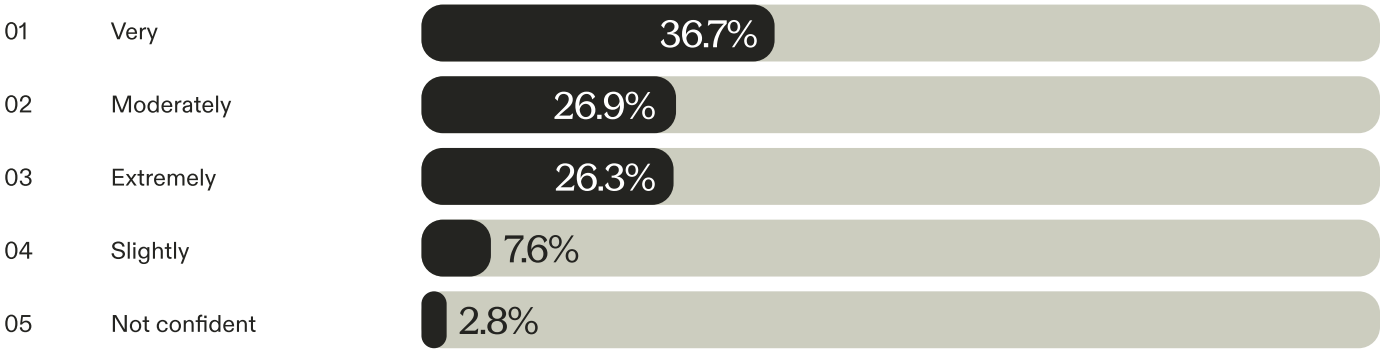
AI is making the work lighter before it's making the outcomes better, and that shows up clearly in where teams report the most improvement: efficiency, lower costs, and productivity gains.

Where teams have seen measurable improvment



Teams that are seeing value today are confident it'll keep paying off. **Nearly 90% feel at least moderately confident that AI will continue to show value over the next year.** Just 2.8% aren't confident at all.

Confidence in agentic AI's value over the next 12 months



Nobody's questioning AI's place on the team, just its answers and tools around it.

4.4%

of teams fully trust AI-generated outputs without review

Almost nobody takes AI's output at face value. More than half of respondents reported having to do at least some checking, and more than a third say outputs still require heavy human review before they'll rely on what it produces.

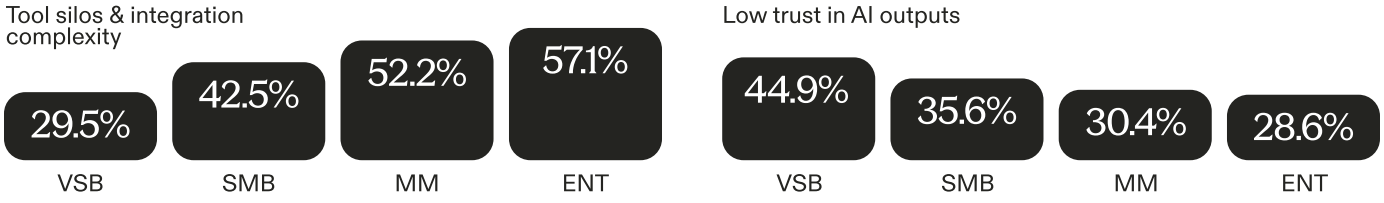
4.4% of teams fully trust AI-generated outputs without review. That's a hard number for any vendor pitching full autonomy as the end state. What's actually working on the ground is a split: automate the repetitive steps, keep a human reviewing anything that touches a prospect or a deal. Teams have already built that balance into their workflows, well ahead of any vendor's autonomy roadmap.

And to be clear, none of this is about whether AI belongs on the team. It's already there. This is about whether they can trust what it just told them.

That trust gap isn't the same for every team, though, and it's not the only barrier standing in the way.

Team size decides whether AI's biggest barrier is trust or tooling

Bigger teams are less likely to struggle with trust. What trips them up instead is infrastructure, tool silos, and integration complexity. Smaller teams face the opposite mix: simpler tooling, but a harder time trusting AI enough to rely on it.



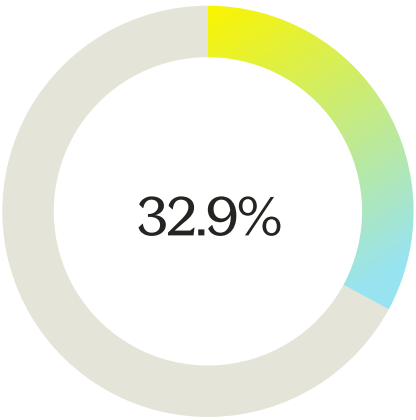
The future of GTM is architecture, not add-ons.

74%

[of teams] report running somewhere between 2 – 5 different GTM tools today

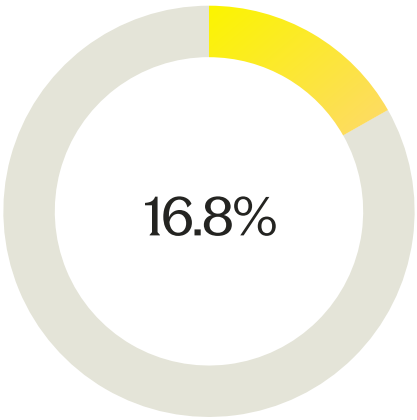
Most teams already have plenty of tools. **74% report running somewhere between two and five different GTM tools today.** With that much already in place, the real question for 2026 isn't which tool to add next. It's what to do with the ones already there. Automating the workflow end-to-end is the clear #1 answer, nearly double the next closest priority.

GTM teams' #1 priority for agentic AI in 2026



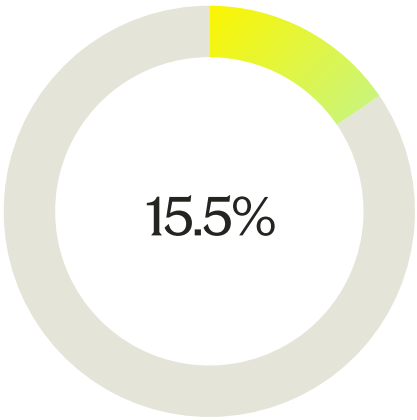
01

Automate workflows end-to-end



02

Consolidate tools into a unified workflow



03

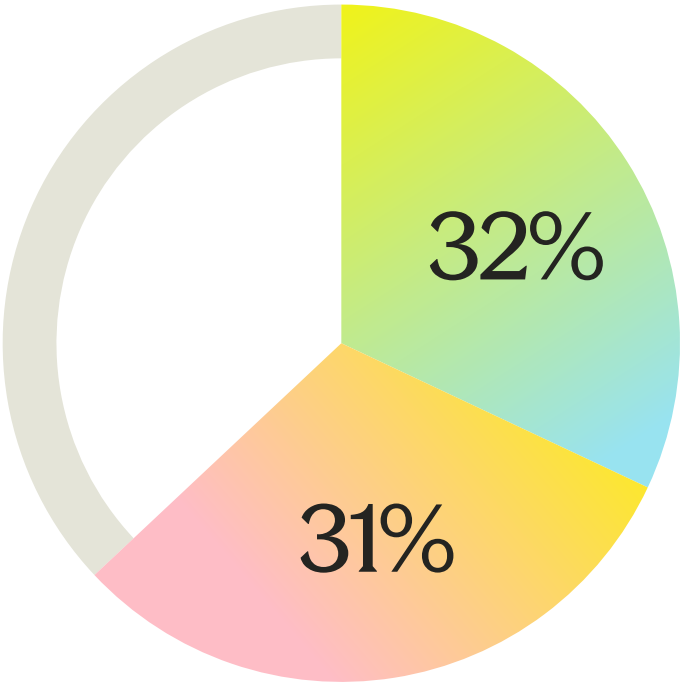
Prove ROI with better measurement

If the goal is one connected motion, the obvious next question is where that motion actually starts. Right now, it starts in two different places, almost evenly split.

About a third of teams start in a SaaS tool, like a CRM, and bring in an LLM to refine the work. Nearly as many start the other way, opening an LLM first and pushing the output into their SaaS stack afterward.

Where the motion starts

Near even split is a sign of friction.



- Start in a SaaS tool, then bring in an LLM to refine the work **32%**
- Start in an LLM, then push outputs into SaaS tools **31%**
- Other..... **37%**

That near-even split is a sign of friction. Teams are manually stitching two disconnected surfaces together (copying context from a CRM into a chat window, then pasting the output back into the CRM). Leaders are aware though, and that's why automating the workflow end-to-end is the top priority for 2026.

Whichever surface a team starts on, the outcome they're after is the same: one connected motion, not a choice between tools.

Four steps to close the AI gap

01	Define your terms	"Agentic AI" means different things to different teams. Align on a definition before you align on a strategy.
02	Connect before you collect	Most teams already have 2-5 GTM tools. The priority for 2026 is making them work together—or consolidating to fewer, more capable tools—not adding more.
03	Build trust through checking, not blind faith	Almost nobody trusts AI output without review. Build a review habit instead of avoiding AI or over-relying on it.
04	Start wherever you already are	Whether that's a SaaS tool or an LLM, they should end up working as one connected GTM system, not siloed tools.

SPOTLIGHT

Apollo, wherever you work.

Apollo is the AI GTM System that shows up wherever a team already works. Start in Apollo, your CRM, or in whatever AI tool's already open: ChatGPT, Claude, Perplexity, a custom build. Then bring Apollo's data and execution with you, through the API, CLI, and MCP.

Learn more

Customer testimonials



Sam Knollmeyer
Strategy & Ops
Alternative Payments

Alternative Payments

"We do pretty much everything in Apollo by API. Their CLI + MCP have become powerful enough that we consolidated our stack, and now use Apollo's headless product for core enrichment while we build custom scoring, routing, and signal workflows on top of it."



Marc Daniels
Head of GTM
Magenta

MAGENTA

"We've used Apollo MCP to scale outbound operations and integrate data with our proprietary sources. This wouldn't have been possible if Apollo hadn't invested in tools compatible with modern AI models like ChatGPT, Gemini, and Claude."

Ready for more?

Sign up for free

Get a demo