

PROMPT UNLOCKED | FREE MONTHLY REPORT

The State of AI Tools & Agents

September 2026

Adoption stats, market share and what it means for prompt writers. Published monthly by promptunlocked.com

1. Adoption: The Two-Track Gap

THE HEADLINE NUMBERS

McKinsey's State of AI global survey, fielded 4 May to 8 June 2026 across 1,719 respondents in 97 nations, is the most current broad dataset available this month. It shows adoption widening and deepening at the same time.

44%

of organisations now report AI scaling across the enterprise, up from 38% a year ago

56%

use AI in three or more business functions, up from 51%

54%

of organisations above \$1bn revenue are scaling AI, against roughly a third of smaller ones

The pattern worth noticing is not the growth itself but the widening split by company size. Larger organisations have moved faster than smaller ones over the past year specifically on agentic AI, and secondary analyses of the same survey put the share of \$1bn-plus companies scaling agents at 40%, up from 27%, while smaller firms stayed roughly flat.

THE GAP NOBODY HAS CLOSED

Conviction is running ahead of measurable return. Around 80% of respondents report individual productivity gains, but only about 37% see organisational financial impact, and the share of genuine high performers has not moved year on year. McKinsey's own framing is blunt: conviction in AI is growing faster than the financial returns organisations can attribute to it.

For anyone writing prompts or building workflows, that gap is the opportunity. The bottleneck is no longer access to capable models. It is the translation of individual speed into something that shows up in the accounts.

2. Agents: Experiment vs Production

WHERE AGENTS ACTUALLY ARE

Agent adoption reads very differently depending on which denominator you use, and most headlines pick the flattering one. Roughly two thirds of enterprises have experimented with agents. Under 10% have scaled them in any single business function. Gartner's CIO survey puts deployed-agent adoption at around 17%. All of these can be true at once.

The most useful number this month is a different one: McKinsey found that 32% of organisations skipped buying at least one software product or feature because they could build it internally with agentic coding tools. That is agents changing procurement decisions, not just workflows.

WHAT LAUNCHED THIS MONTH

- Grok Bot (xAI, beta since 11 August 2026) gives each assigned task its own cloud computer, its own logins and a runtime that persists after you close your laptop. Memory survives restarts.
- Claudeforce, a Salesforce and Anthropic partnership, embeds Claude as a default reasoning model across Agentforce.
- Cisco extended its MyAgent programme to roughly 90,000 employees, a rare example of company-wide internal agent deployment rather than a pilot.

The common thread is persistence and scope. The interesting agent products this month are not better chatbots; they are processes that keep running when you are not watching. That changes what can go wrong.

3. What It Means for Prompt Writers

FOUR PRACTICAL SHIFTS

1. Briefing beats prompting. When an agent runs unattended, the scarce skill stops being clever phrasing and becomes precise scoping: what success looks like, what the agent must never touch, and where it stops to ask. A vague prompt returns a bad answer; a vague brief runs all night.
2. Failure modes changed. A chat gives a wrong answer you can see. An agent with credentials and a persistent runtime acts on the wrong answer repeatedly. Caps, approval gates for anything irreversible and full execution logs are no longer optional extras.
3. The build-vs-buy line moved. With a third of organisations already skipping software purchases because agentic coding tools let them build instead, being able to specify an internal tool well is worth more than knowing which SaaS product to buy.
4. Structure is reliable, facts are not. This holds across every category we track. Generated decks, articles and reports are consistently well-organised and consistently confident about specific numbers that are wrong. Verification is the job that did not get automated.

SOURCES AND CAVEATS

Primary source: McKinsey, The State of AI global survey, fielded 4 May to 8 June 2026, 1,719 respondents across 97 nations. Agent deployment figures are drawn from Gartner CIO survey data and secondary analyses of the same McKinsey dataset, which is why they differ from one another. Read every adoption statistic by asking what is being measured, who was surveyed, and whether the number reflects usage, rollout or business outcome.

Product launch dates are from vendor announcements. Pricing changes frequently and is deliberately omitted here; check the vendor before budgeting. This report is free, independent and not sponsored by any tool mentioned.

promptunlocked.com | Next edition: October 2026