



2026 Developer's Guide to **Agentic AI**

What developers and product leaders are actually building,
trusting, and shipping today

nylas.com/agicntic-ai-report-2026

Table of contents

Research methodology	3
The builders shaping agentic AI	4
Agentic AI 2026 snapshot	5
No one agrees on what agentic AI means. Not even developers	6
Agentic AI is already shipping	7
Where agentic AI is already proving value	9
Autonomy is not all or nothing	10
Agentic AI will be table stakes within three years	12
Why teams aren't panicking about competitors	13
From side project to roadmap	14
Speed beats cost (for now)	15
What's holding teams back	16
Agentic AI is a 'switching event'	17
Conclusion: Agentic AI will be defined by what ships	18
FAQ	19
Glossary	20

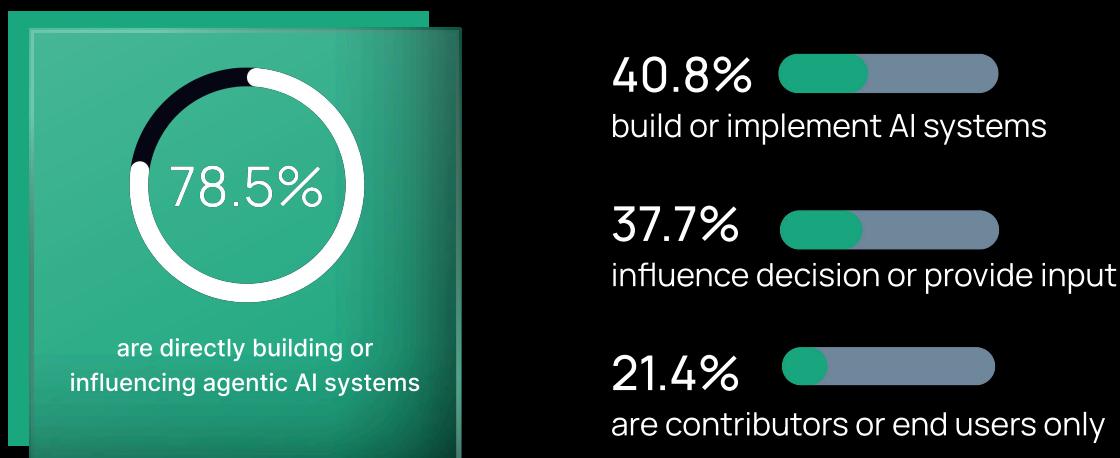
Research methodology

The findings in this report are based on an online survey of **1,026 respondents** conducted between December 18–30, 2025 with the assistance of [Centiment](#). Respondents were primarily based in the United States and work across engineering, product management, IT, data, and executive leadership roles.

Nearly **80% of respondents are directly involved in building or influencing AI and agentic AI decisions** at their organizations. Most respondents work at mid-to-large companies, with a strong concentration in software and SaaS, as well as meaningful representation from finance, healthcare, consulting, and customer experience teams.

Respondents were not selected based on usage of any specific vendor or platform. This report reflects the perspectives of practitioners actively shaping how agentic AI is designed, trusted, and deployed in real production environments. Given the rapid pace of tooling innovation in agentic AI, particularly around autonomous execution, adoption patterns may continue to evolve quickly beyond this snapshot in time.

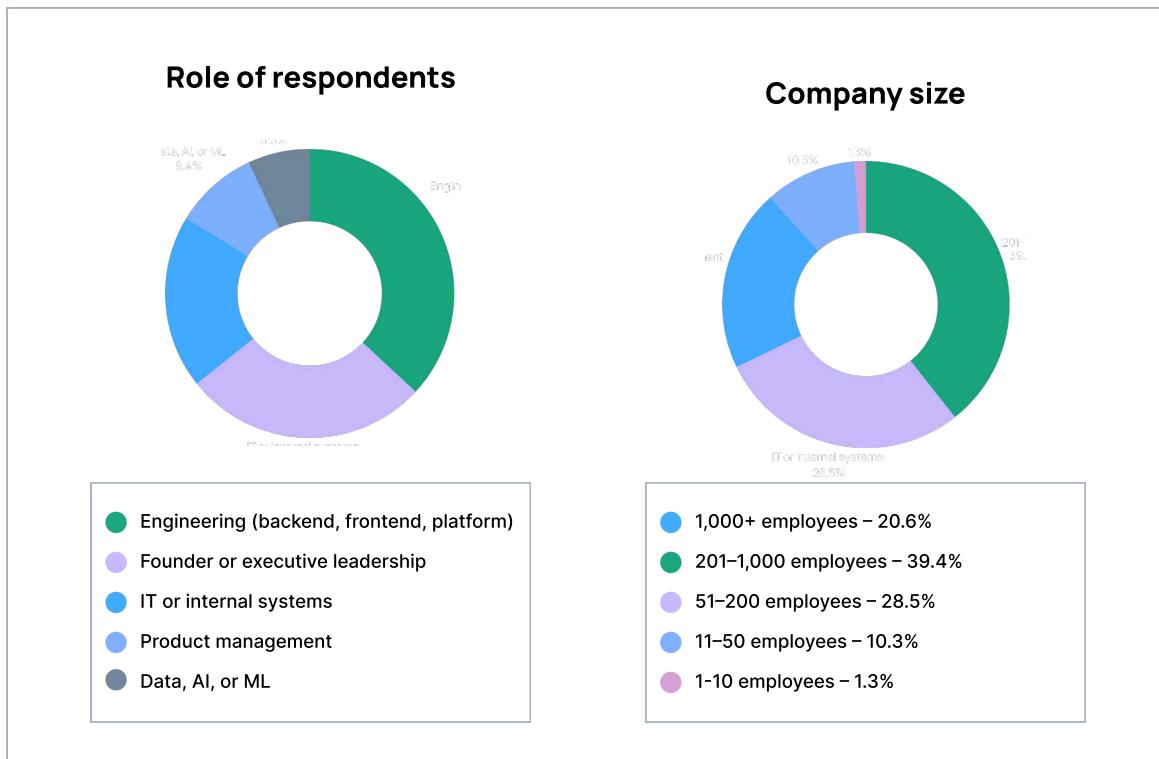
How involved respondents are in making AI decisions



The builders shaping agentic AI

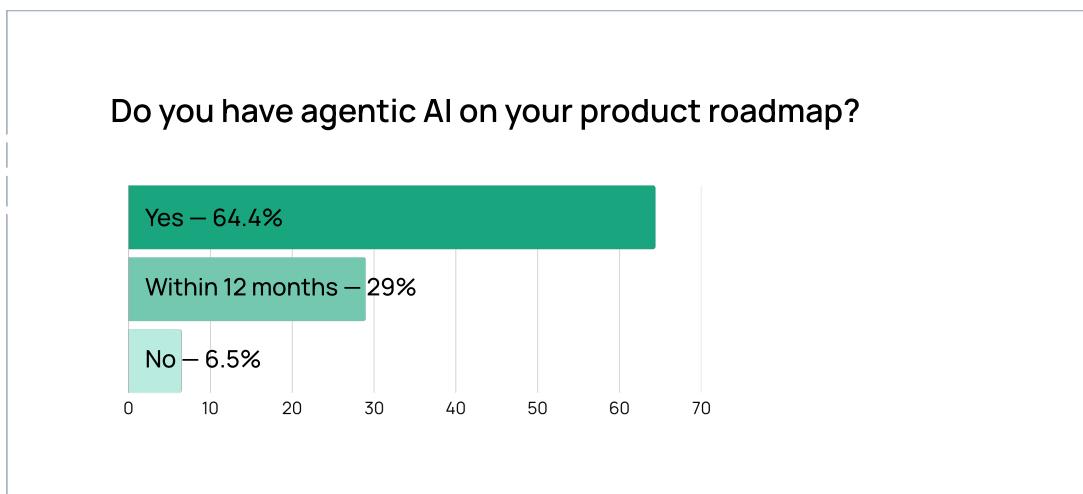
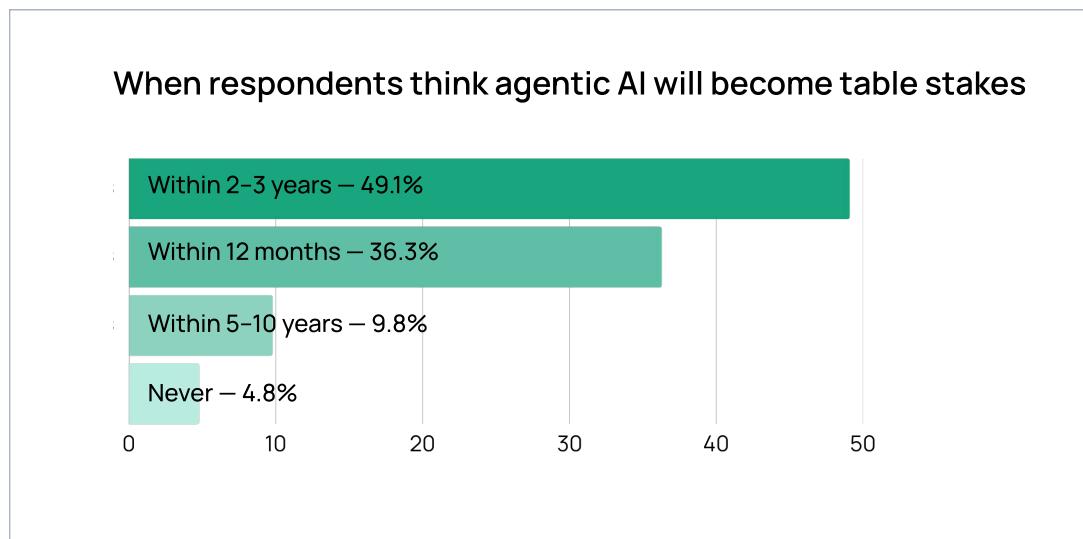
Agentic AI is being shaped by builders, not spectators. The majority of respondents work in engineering, product management, or IT and internal systems roles. Most are employed at organizations with more than 200 employees, and many work at companies with over 1,000 employees.

This is not an audience experimenting on the margins. These are the teams responsible for production systems, infrastructure decisions, security approvals, and long-term roadmaps. What they build and trust today is what ends up defining agentic AI in practice.



Agentic AI 2026 snapshot

These are three signals we couldn't ignore. Agentic AI is becoming table stakes faster than expected. It's already on most product roadmaps and strong enough to drive real vendor switching. This page offers a snapshot of our findings. The full Developer's Guide to Agentic AI 2026 goes deeper with 20+ charts on how teams are building, governing, and scaling agentic systems today.



Likelihood to switch vendors for stronger agentic AI capabilities

Very likely

Possibly

Unlikely

No

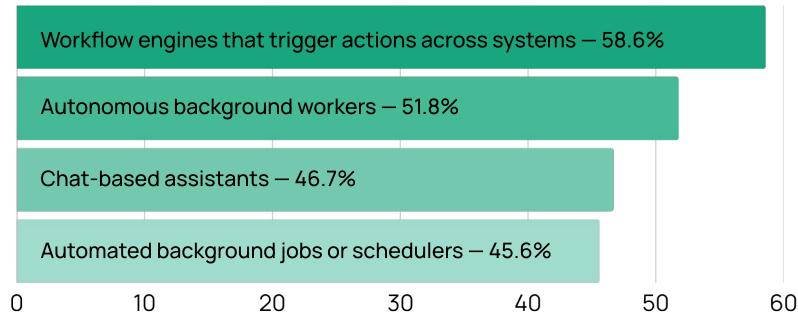


No one agrees on what agentic AI means. Not even developers.

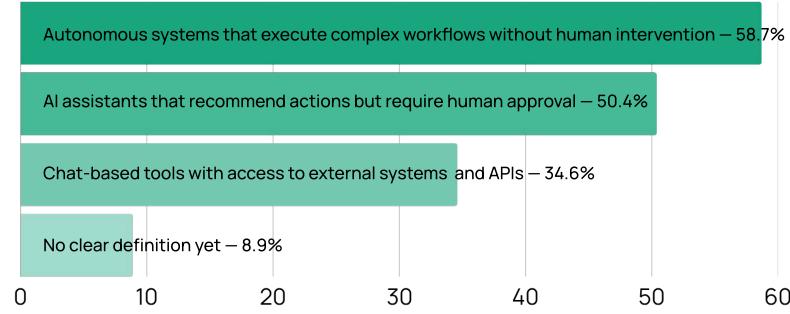
Ask ten teams what “agentic AI” means, and you will get ten different answers. There is no shared definition today. Some teams use the term to describe workflow engines that trigger actions across systems. Others mean autonomous background workers, human-in-the-loop assistants, or chat-based tools with access to APIs.

The data backs this up. When asked what agentic AI refers to, respondents split across multiple interpretations, with no single model clearly leading. Even at the organizational level, teams remain divided between fully autonomous systems, assisted workflows, and chat-based tools connected to external systems.

What do you think agentic AI means?



Which best describes your current definition of agentic AI?



That confusion is not a failure. It is a signal. This category is still being shaped in real time by what teams are building, testing, and slowly learning to trust in production. Agentic AI will not be defined by a whitepaper or a glossary. It will be defined by the first workflows teams cannot imagine turning off.

Research findings

Agentic AI is already shipping

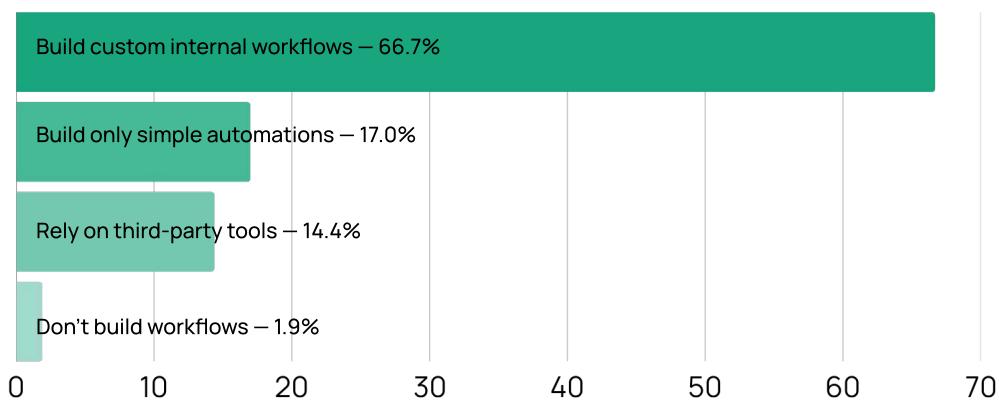
Agentic AI is already showing up in real systems, even if you don't always see it from the outside. **67% of respondents are building custom internal workflows**, and nearly half say they are actively building agentic AI today. This isn't future work. It's already in production.

Definitions are all over the place, but behavior isn't. Most teams are already building something agentic, whether they call it that or not. Two-thirds report building custom internal workflows today. Nearly half are actively building agentic AI, with another third seriously evaluating it. Very few teams are just watching.

Most of this starts internally. Engineering automation. IT workflows. Internal services. Operational tooling. Not splashy product launches. These are the places teams can move fast, try things, and learn what breaks without a lot of risk.

That gap matters. Public conversations make agentic AI sound new and experimental. Inside companies, it already feels normal. Teams are putting agents to work, seeing how they behave, and adjusting as they go. By the time this shows up as a headline feature, a lot of teams have already been living with it for a while.

How is your team building internal automation or workflows today?

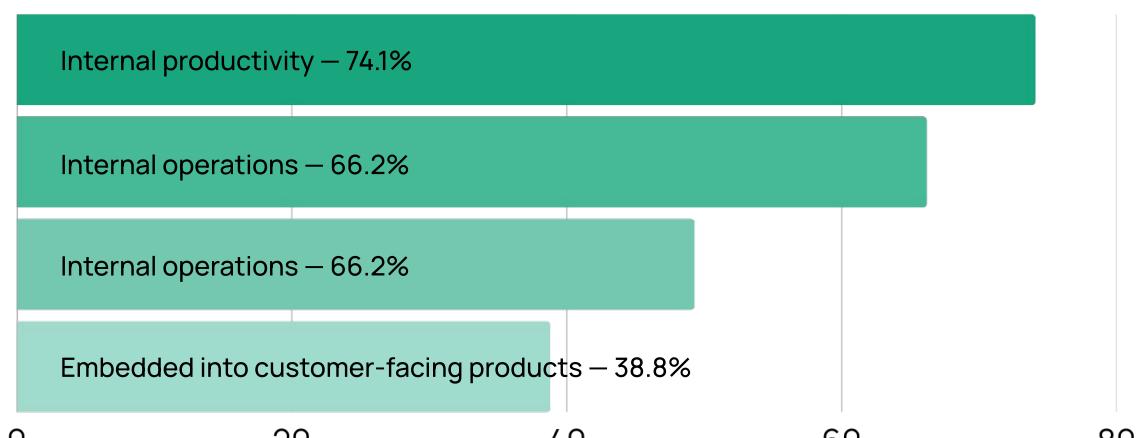


If agentic AI is already shipping,
why does it still feel experimental?

Which best describes your relationship with agentic AI right now?



How do you use AI tools today?



Internal productivity and operations lead agentic AI adoption, with customer-facing use cases still emerging

Where agentic AI is already proving value

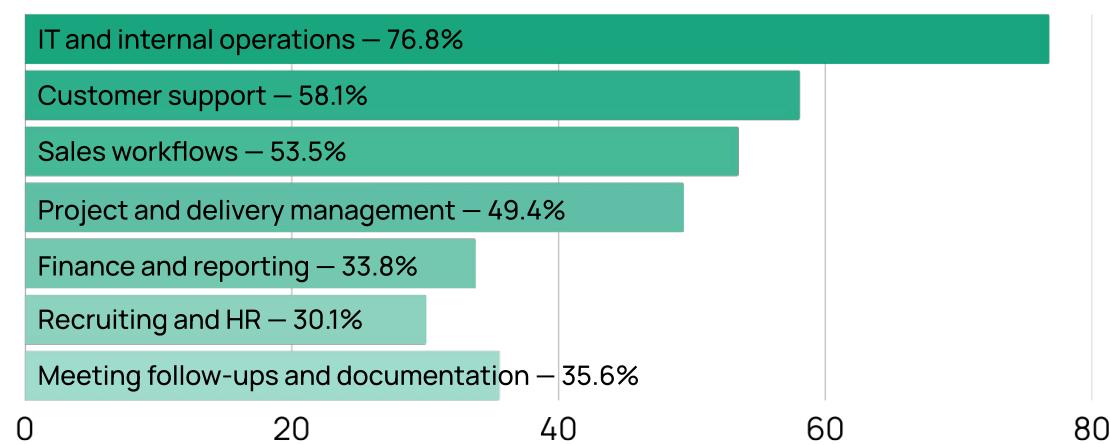
Agentic AI shows up first where speed and coordination actually matter. Not everywhere at once. In very specific workflows where things are repetitive, time-sensitive, and spread across too many systems.

IT and internal operations lead the way, followed closely by customer support, sales workflows, and project and delivery management. These are the places where handoffs pile up, context gets lost, and small delays turn into real pain. That's where agents start to earn their keep.

More visible use cases are still catching up. Meeting follow-ups and documentation are starting to show traction, but they haven't reached the same level of maturity yet. Most teams are still focused on getting the internal stuff right first.

That pattern makes sense. Agentic AI gains traction where it can quietly remove friction from complex, repetitive work before it ever shows up as a polished, customer-facing feature.

Which workflows have you seen agentic AI influence?



Agentic AI is delivering the most value today in internal and operational workflows where coordination and speed are critical.

Does trusting agents mean giving up control?

Autonomy is not all or nothing

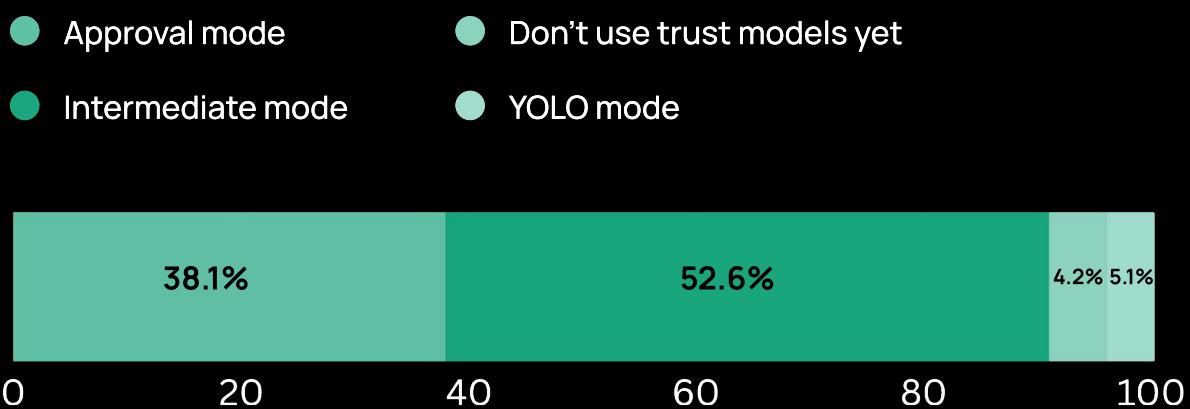
At the time of the survey, only **4%** of teams allow agents to act without human approval. Full autonomy is still rare. That said, it's clearly picking up speed as tooling gets better and teams get more comfortable letting agents take on real work.

Most teams aren't flipping a switch. They're easing into it. Agents handle the routine, low-risk stuff. Humans stay in the loop for decisions that actually matter. It's not all-or-nothing autonomy. It's trust built a little at a time, in production.

Public conversations about agentic AI tend to frame autonomy as a binary choice. Either you trust agents or you don't. In practice, that's not how it works. Teams are choosing something far more nuanced.

Most organizations rely on graduated trust models. Some actions are automatic. Others still need approval. Fully autonomous systems without oversight are the exception, not the rule. Only four percent of respondents say they're running agents with no human approval at all. Most prefer an in-between setup that balances speed with control.

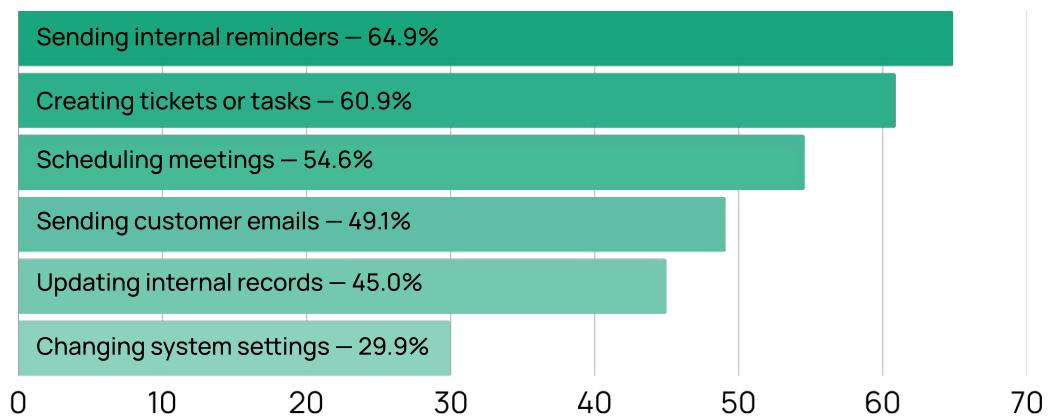
How are you building trust with your AI tools today?



At the same time, trust is expanding. Teams are already comfortable letting agents create tickets, send reminders, schedule meetings, update records, and even send customer emails without approval. These aren't toy examples. They're real workflows that save time and remove friction.

Agentic AI adoption isn't about giving up control. It's about delegating responsibility carefully, step by step, and seeing what holds up.

Which tasks would you trust an agent to perform without approval?



Teams are most comfortable delegating routine, low-risk actions to agents without human review.

Agentic AI will be table stakes within three years

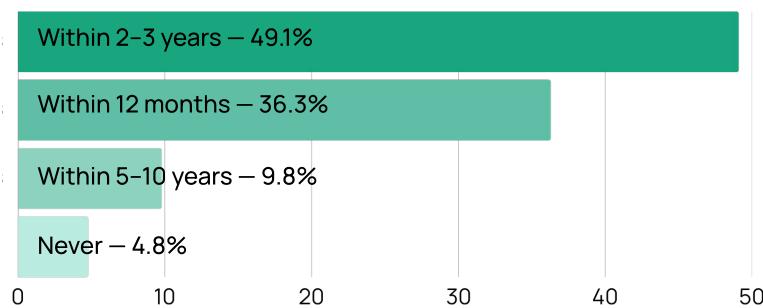
Teams don't agree on what agentic AI should look like, but they do agree on one thing: it won't be optional for long. **85% believe agentic AI will become table stakes within the next three years**, and more than a third think that happens within the next 12 months.

That's not some distant horizon. It's close enough that teams are already planning around it, even while they're still arguing about definitions. Agentic AI is moving out of the "nice-to-have experiment" bucket and into the category of things teams expect to be built in.

Only a small minority think agentic AI won't become essential. Everyone else sees the same trajectory, even if they disagree on the exact shape it'll take. The debate isn't really about if this matters anymore. For many teams, that work is starting internally, but it still requires the same infrastructure, guardrails, and best practices they'll eventually need in customer-facing products. It's about how fast teams can make it work without breaking trust, security, or control.



When will agentic AI become table stakes?



Most teams expect agentic AI to become essential within the next three years.

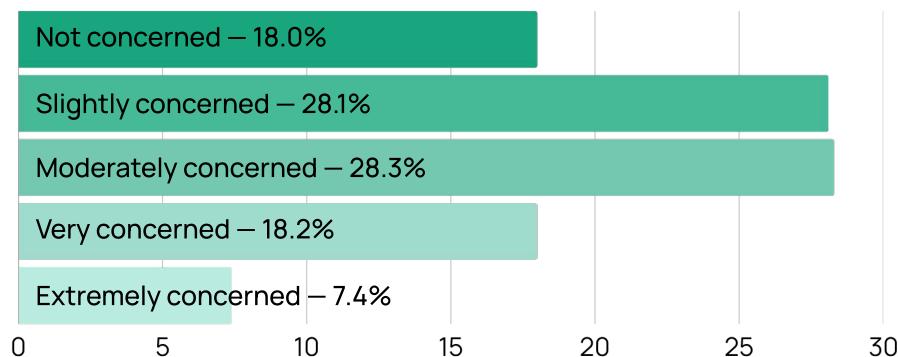
Why teams aren't panicking about competitors

Even with how fast things are moving, most teams aren't panicking about competitors shipping agentic AI before them. The concern is there, but it's muted. If we can't decide on what it even is, why should we care if someone else ships it? Most respondents fall into the "slightly" or "moderately" concerned range. Only about a quarter say they're very or extremely concerned.

That posture makes sense. Many teams are already building internally, and they feel early enough to keep pace. They're not standing still. They're learning by doing, even if what they're building isn't public yet.

That confidence isn't denial, though. It's conditional. As internal tools harden and agentic AI starts showing up in customer-facing products, the pressure will rise fast. The window where teams can quietly catch up won't stay open forever.

How concerned are you about competitors shipping agentic AI before you?



Most teams report moderate concern, with relatively few expressing high levels of urgency.

From side project to roadmap

Agentic AI isn't sitting on the sidelines anymore. Nearly two-thirds of respondents already have it on their product roadmap. More than half say it's critical or very important to their overall strategy. Very few see it as overhyped or irrelevant.

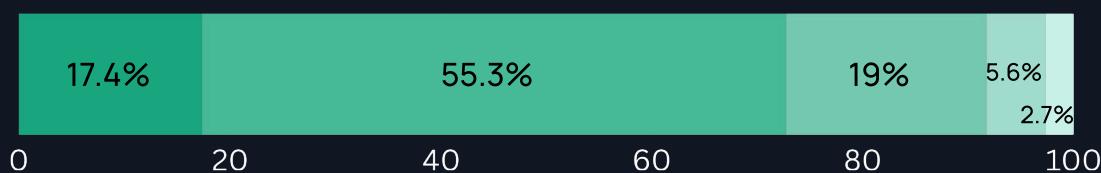
That's a real shift. This isn't just something teams are poking at on the side. It's showing up in planning conversations, roadmap reviews, and longer-term bets.

Once something makes it onto the roadmap, it starts to change how teams think. It affects where they invest, what they prioritize, and which infrastructure decisions they're willing to make. Agentic AI is moving out of the "let's try this" phase and into the way products actually get built.



How important is agentic AI to your overall product strategy?

- Critical
- Very important
- Somewhat important
- Low priority
- Not relevant



For most teams, agentic AI has moved from experimentation into core roadmap planning.

If speed matters so much,
what's slowing teams down?

Speed beats cost (for now)

Teams aren't adopting agentic AI to save a few dollars. They're doing it to move faster. Improving response time ranks ahead of cost reduction as the top reason teams are investing here. Speed is the unlock.

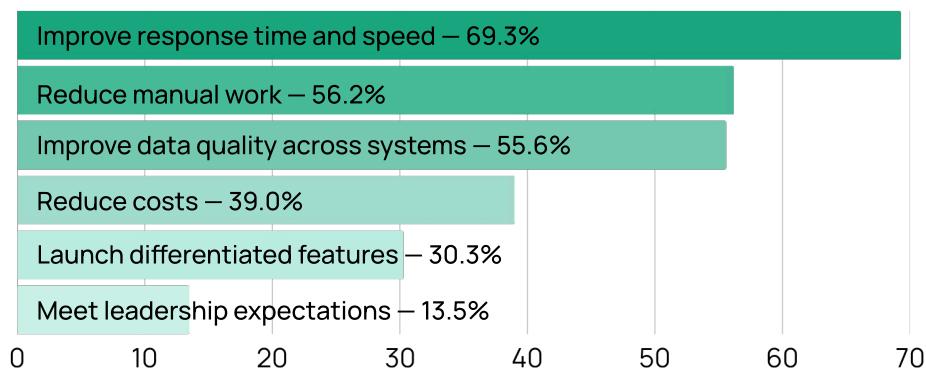
But speed on its own isn't enough. Once agents start touching real workflows, reliability, security, and integration coverage matter just as much. If something breaks, misfires, or doesn't connect cleanly to the rest of the stack, the speed gains disappear fast.

That's changing how teams evaluate tools. Feature lists matter less than whether something actually holds up in production. Can it run every day without babysitting? Can it plug into the systems teams already rely on? Can it be trusted when things go sideways?

When teams talk about why they're adopting agentic AI, the pattern is clear. Speed comes first. Reducing manual work and improving data quality follow. Cost savings trail behind. For most teams, this isn't about operating cheaper. It's about not falling behind.

Agentic AI isn't a cost-cutting tool. It's a velocity tool. And every company without an agentic workflow just got slower.

Why teams are adopting agentic AI



Speed and responsiveness are the primary drivers of adoption, outweighing cost reduction or executive pressure.

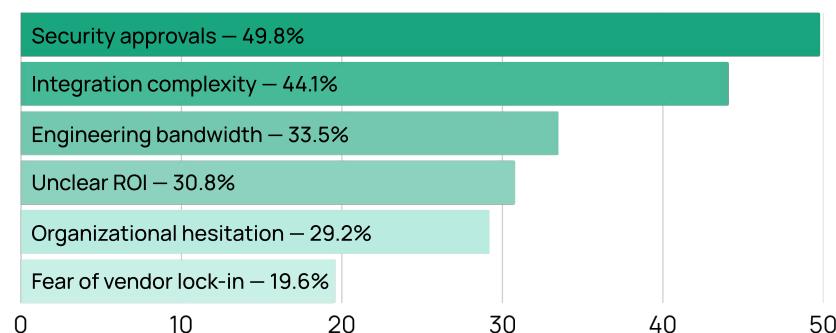
What's holding teams back

Momentum isn't the problem. Constraints are. Teams generally know what they want to build with agentic AI. What slows them down are the practical blockers. Security approvals. Integration complexity. Limited engineering bandwidth. Unclear ROI. These aren't philosophical debates. They're day-to-day constraints around permissions, systems, and coordination.

On the technical side, the hard parts show up fast. Reliability. Failure handling. Permissions and identity. Compliance. Cross-platform integration. These are the things that make or break agents once they move past a demo. Platform changes and API updates only add to the friction, introducing instability that can stall development entirely.

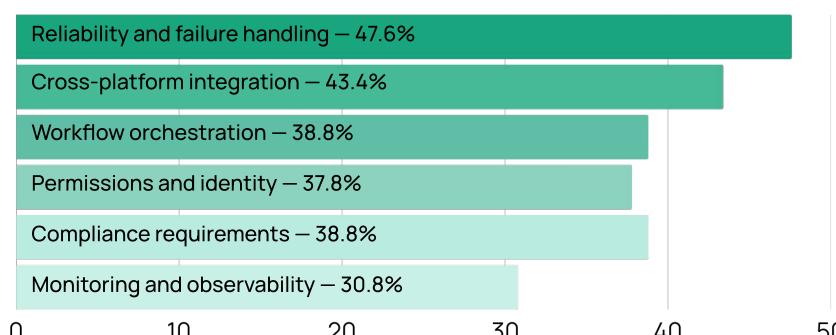
None of this is about a lack of ambition. The biggest barriers to agentic AI adoption are infrastructure, integration, and trust. As agentic AI moves from internal tooling into more visible, customer-facing experiences, those constraints start to matter even more. They don't just slow teams down. They shape which vendors teams stick with and which ones they walk away from.

What's preventing teams from implementing agentic AI



Speed and responsiveness are the primary drivers of adoption

The hardest technical problems teams face when building agents



Reliability, integration, and permissions emerge as the most difficult challenges in production environments.

Agentic AI is a 'switching event'

When does agentic AI stop being a feature and start being a reason to switch platforms?

Agentic AI is already changing how teams think about buying tools. **94% of respondents say they'd possibly or very likely switch vendors for stronger agentic AI capabilities.** That's not a gentle preference shift. That's a real willingness to move.

This pushes agentic AI out of the "nice feature" bucket. It starts to act like a forcing function. Teams are deciding which platforms are worth committing to, and which ones aren't keeping up.

What's driving those decisions isn't novelty or flash. It's the basics, done well. Reliability. Integration coverage. Security. Teams are looking for agentic systems they can actually trust to run inside real workflows, not just demo well. If an agent is going to touch production systems, it has to work every time and play nicely with everything else.

The result is churn pressure. And consolidation. And displacement. Agentic AI is becoming an acquisition and retention lever whether vendors planned for it or not. Teams are voting with their feet, and platforms that can't deliver reliable, well-integrated agentic experiences are increasingly at risk of being replaced.

Likelihood to switch vendors for stronger agentic AI capabilities

● Very likely ● Possibly ● Unlikely ● No



Most teams are open to switching tools or platforms if stronger agentic AI capabilities are available.

Conclusion: Agentic AI will be defined by what ships

Agentic AI won't be defined by definitions, frameworks, or hype cycles. It won't be settled by blog posts or positioning decks. It'll be defined by what actually survives contact with real users.

The winners won't be the teams with the cleanest theory. They'll be the ones that ship workflows that hold up. Systems that work every day. Infrastructure teams trust with real data, real permissions, and real consequences.

The people building agentic AI today aren't waiting for consensus. They're experimenting, shipping, breaking things, fixing them, and learning in parallel. Most of this work is happening quietly, inside internal tools and operational systems, long before it shows up as a headline feature.

The category will solidify when one of those experiments becomes something teams can't imagine turning off. When it stops feeling like "agentic AI" and starts feeling like part of how work gets done.

The teams that ship first won't just win the category. They'll define what agentic AI even means.

The teams that **ship first won't just win the category.
They'll **define what agentic AI even means.****

FAQ about agentic AI

What is agentic AI?

Agentic AI is a class of AI systems that can plan, decide, and take action across workflows with varying levels of autonomy. Unlike traditional AI models that generate single responses, agentic AI systems can execute multi-step tasks, use tools and APIs, maintain context, and adapt based on outcomes.

In practice, agentic AI is used to automate workflows, coordinate actions across systems, and handle tasks that would otherwise require human intervention. Most real-world agentic AI systems operate with guardrails and human oversight rather than full autonomy.

What's the difference between agentic AI and AI agents?

AI agents are individual software entities that can observe context, make decisions, and take actions, often by calling tools or APIs. An AI agent is a building block.

Agentic AI refers to the broader system or capability that emerges when one or more AI agents are used to execute workflows with autonomy, coordination, and oversight. In short, AI agents are the components. Agentic AI is how those components are designed, trusted, and deployed in real systems.

What makes an AI system “agentic”?

An AI system is considered agentic when it can plan and execute multi-step tasks, take real actions using tools or APIs, maintain state or memory across steps, and operate with some level of autonomy.

Most agentic systems also include safeguards, such as human-in-the-loop approval or graduated trust models, especially for higher-risk actions.

Is agentic AI fully autonomous?

Most agentic AI systems are not fully autonomous. In production environments, teams typically use graduated trust or human-in-the-loop models where low-risk actions are automated and higher-risk actions require approval.

Fully autonomous agents without oversight remain rare and are usually limited to narrow, low-impact workflows.

What is agentic AI used for today?

Agentic AI is most commonly used for internal productivity and operations. Common use cases include IT workflows, internal automation, customer support coordination, sales workflows, scheduling, follow-ups, and data updates across systems.

Many teams start with internal use cases before expanding agentic AI into customer-facing products.

Glossary

Agentic AI: AI systems that can plan, decide, and take action across workflows with varying degrees of autonomy. Agentic AI typically combines reasoning, memory, and tool use to execute multi-step tasks, not just generate single responses. (See also: AI agent, Autonomy, Workflow orchestration)

Agentic AI vs AI agents: Agentic AI refers to the broader category or capability set, while AI agents are the individual components that carry out actions. In practice, agentic AI systems are made up of one or more AI agents working together across workflows.

Agentic workflow: A multi-step process executed by one or more AI agents across systems, with varying levels of autonomy and human oversight. Agentic workflows often span multiple tools and data sources.

AI agent: A software entity powered by AI that can observe context, make decisions, and perform actions on behalf of a user or system. AI agents often interact with APIs, databases, and other tools to complete tasks end to end. (See also: Agentic AI, Tool calling)

Autonomy: The degree to which an AI system or agent can act without human intervention. Autonomy exists on a spectrum, from fully human-approved actions to fully autonomous execution. (See also: Human-in-the-loop, Graduated trust model)

Customer-facing agents: AI agents embedded into user-facing products or experiences. These agents interact directly with end users and typically require higher levels of trust, security, and reliability.

Failure handling: Mechanisms that allow AI systems to recover from errors, partial failures, or unexpected states. This includes retries, rollbacks, alerts, and safe fallbacks when actions cannot be completed.

Graduated trust model: An approach to agentic AI where autonomy is granted incrementally. Agents are allowed to perform low-risk actions automatically, while higher-risk actions require human approval or additional safeguards.

Human-in-the-loop (HITL): A trust and control model where AI systems require human review or approval for some or all actions. Human-in-the-loop approaches are commonly used for higher-risk decisions, external communications, or sensitive data access. (See also: Autonomy, Graduated trust model)

Integration complexity: Challenges related to connecting AI systems across multiple platforms, APIs, identity systems, and data sources. Integration complexity increases as workflows span more tools and require consistent permissions and reliability.

Glossary, Continued

Internal agents: AI agents used inside an organization for internal productivity, operations, IT, or engineering workflows. Internal agents are often where teams first experiment with and validate agentic AI. (See also: Customer-facing agents)

Observability: The ability to monitor, debug, and audit AI system behavior. Observability includes logging decisions, tracking actions taken by agents, and understanding why a system behaved a certain way.

Permissions and identity: Controls that determine what actions an AI system or agent is allowed to perform and on whose behalf. Permissions and identity management are critical for securing agentic workflows across systems. (See also: Security and compliance)

Reliability: The ability of an AI system or agent to perform consistently and correctly over time. For agentic AI, reliability includes handling failures gracefully, avoiding duplicate actions, and maintaining correct state across workflows.

Security and compliance: The policies, controls, and safeguards that ensure agentic AI systems operate safely, legally, and in line with organizational requirements. This includes data access, auditability, and regulatory compliance.

Switching event: A moment when a new capability, such as agentic AI, becomes important enough to drive customers to switch tools or vendors. Switching events often reshape markets and accelerate consolidation.

Table stakes: Capabilities that are considered essential rather than differentiating. When agentic AI becomes table stakes, teams expect it to be built into platforms by default.

Tool calling: The ability for an AI system or agent to invoke external tools, APIs, or services as part of a workflow. Tool calling enables agents to take real actions, such as sending emails, updating records, or scheduling meetings.

Workflow orchestration: The coordination of tasks, tools, APIs, and systems to execute multi-step processes. In agentic AI, orchestration allows agents to move between systems, manage state, and complete workflows end to end. (See also: Agentic workflow, Tool calling)



Nylas empowers developers with powerful integrations for email, calendar, and meeting data. We're the only platform that unifies connections to over 250 communications providers, including Google, Microsoft, IMAP, and Zoom, saving your development team months of building infrastructure. We help you unlock valuable insights from communications data faster, so you can ship customer experiences that make an impact in the market.

Visit nylas.com to learn more.