Insight Report™

A Leap of Trust: AI Agents Are Winning Hearts and Wallets

Introduction

AI agents made a great leap across the chasm in 2025. Our data reveals that 3 out of 4 companies have invested in them in the last year, and over half of them are highly likely to expand scope and budget over the next 12 months.

Tim Sanders
Chief Innovation Officer

In contrast to recent reports, such as MIT's The GenAI Divide, that question AI's success in producing value or moving from pilot to production, our findings suggest that agents are succeeding more often than not.

We define agents as systems that can reason, plan, use tools, and, most of all, act. Unlike chat-based generative AI (genAI), which often stays in the background, agents step into the real world: they talk with customers, coordinate with suppliers, and move work through your systems (and with other agents).

In this report, we distinguish two levels of autonomy:

- **Guardrailed autonomy**: Agents work inside agreed boundaries and pause for human approval at key moments
- **Full autonomy**: Trusted agents act first and humans review after, using rollbacks and QA when needed

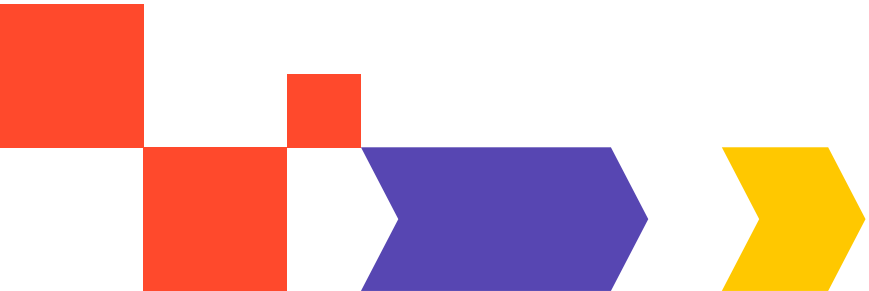

While agents still have a long way to go in terms of reliability and scalability, one thing is clear: they are winning hearts and wallets at companies of all sizes. While many experts argued at the beginning of this year that agents would largely struggle in enterprise workflows, our research tells a different story.

Investor enthusiasm for AI agents is palpable because they translate efficiency into outcomes, something that genAI has yet to deliver consistently, according to many businesses. Fully autonomous agents can remove a primary production constraint: humans managing the loop.

Parkinson's Law, developed by Naval historian Cyril Northcote Parkinson, posits that with humans, the work will always expand to the time afforded. In other words, deadlines drive results, not efficiency gains. Time saved often leads to time wasted instead of bonus production. Studies from academia to software development support this hypothesis.

That's why there's been a disconnect between genAI's productivity improvements and organizational value. Even if humans are faster with a chatbot, that doesn't mean they produce more outputs. Think of this as productivity leakage.

Agents, on the other hand, have the capability to defy Parkinson's Law. They don't operate under deadlines. As they get faster, deliverables roll out quicker. They don't watch TikTok videos, take days off, or get distracted by Slack messages. They grind away. This gives organizations unbridled responsiveness and, from a production standpoint, ramped velocity.

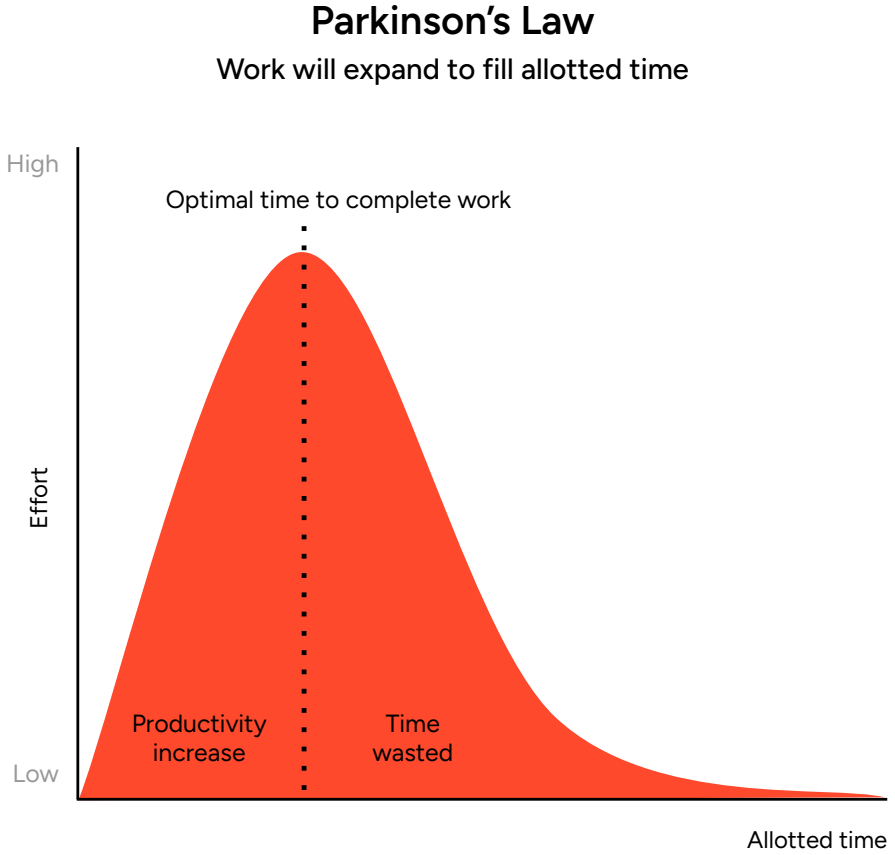

An agent's efficiency is statistically correlated to responsiveness and speed to market. And that changes everything. As researcher and author Jason Jennings wrote,

"These days, it's not the big that eat the small...it's the FAST that eat the slow!"

When companies can trust agents enough to grant them autonomy over critical workflows, they can develop a brute force advantage through a limitless digital labor pool. They will constantly run ahead of schedule as technology improves, leapfrogging those who insist on a human-reliant system of work.

But the secret to giving agents autonomy is for humans to trust the technology and its underlying data. In other words, trust is the gating factor for creating value with AI agents.

If you don't give agents much agency, keeping them on a short leash is just a semi-autonomous solution. It's only incrementally better than chatbot outputs that need to be approved before they're put into production.

Now, it's one thing to trust a technology enough to get your wallet out, but it's another to trust it enough to let it perform actions on your behalf with your systems, your customers, and your code bases with full autonomy.

Both our survey research and data from G2 Reviews reveal that buyers are increasingly trusting agents to be fully autonomous, but only in low-risk workflows. But as you will see, their trust levels are expanding.

This report provides valuable insights for investors, business leaders, vendors, and AI agent enthusiasts. Our August 2025 survey of 1,000+ B2B software buyers, combined with thousands of G2 Reviews, provides a glimpse into the buyer journey from finding agent solutions to scaling them up inside their organizations.

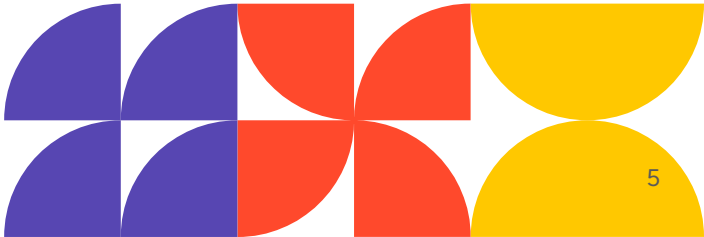

Buyers have told us how the products should function in their environments, what vendors must deliver to earn more trust, and how they prefer to pay for this next-generation AI technology.

We've taken a glimpse into the future with our research, testing a growing belief that “agents will eat SaaS” as an industry. In our view, the transition will produce a new blended landscape instead of a winner-takes-all reality.

One thing is clear:

Five years from now, agents will be leading the charge.

Table of contents

Agent adoption is expanding

Corporate investments in AI agent technology are not trivial. Our survey revealed that 40% of companies have a **$1 million agent budget**, inclusive of software, cloud services, and related staffing. In fact, 1 of 4 large enterprise companies has budgets of $5 million+, outpacing genAI pilot program investments over the last few years.

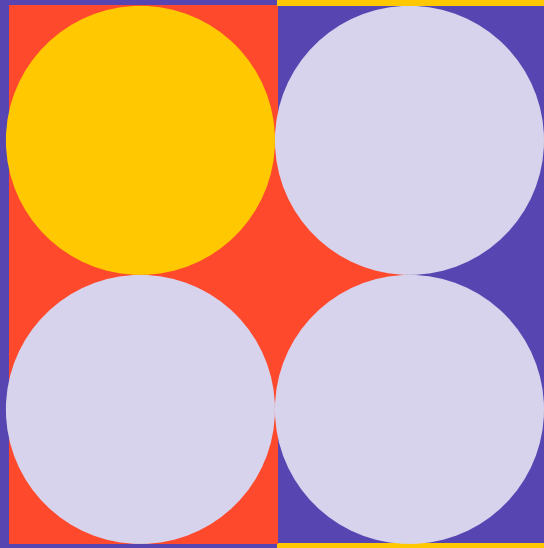

Taking an industry viewpoint, healthcare and retail lead the way in $5 million+ budgets in 2025. However, as this chart illustrates, it's a very nuanced picture. Services, for example, top "no budget" industries, and technology lands somewhere in the middle of the pack.

Further expansion requires companies to move beyond the proof of concept stage to actually deploying agents across key workflows.

Annual budgets for AI agents

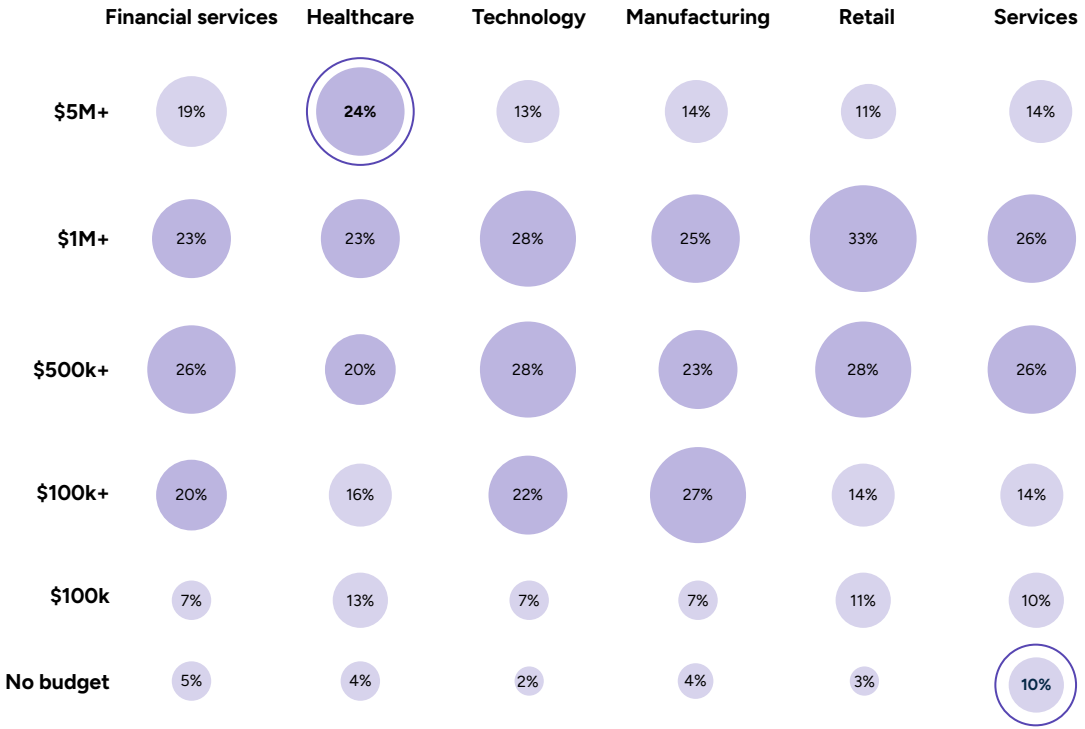

Source: G2 August 2025 Survey

Most B2B companies have AI agents in production

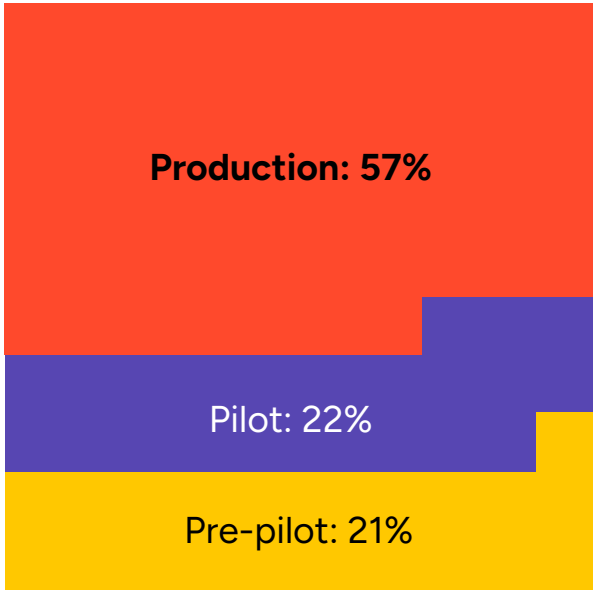

We've measured respondent companies' state of adoption to see how many move beyond the experimental phase to value-creating motions, and the results surprised us. They revealed a very short cycle between test and scale.

Almost 60% of companies have agents right now in production. This is impressive, given how agentic solutions only entered the mainstream conversation last spring.

Our research suggests that before the end of 2026, more than half of the companies that have invested in agents will have them in production, delivering actions and automating business processes.

As noted earlier, recent reports suggest that when it came to genAI, a small percentage of pilots actually made their way into production. Our data suggests that well over half of completed pilots advanced to full production, with a very low percentage failing to make that leap.

Qualitative research compiled from 30+ interviews with technology and IT leaders suggests that the agent buying motion starts with a measurable business problem where the root cause is usually a bottleneck that agents can address.

Source: G2 August 2025 Survey

Agent adoption looks very different from the genAI spending wave of 2023–24. During that period, companies rushed into enterprise-wide seat licenses, carved out budgets for API calls, and even spun up their own models. They went searching for proofs of concept that might show value only after the purchase. The C-suite's fear of missing out to "getting something out of genAI" was at its peak. Today's agent investments follow a different psychology: problem-first, outcome-driven, and rooted in clear ROI.

One way to think about the pilot-to-production process is to look at the most popular use cases for agents, which we've been tracking over the last year.

Top use cases include software development and customer service. Agents are likely trusted here because they squarely address technical and service debt, which is usually due to human capital bottlenecks. Where there is great pain, there is an appetite for solutions, even if they are early in their development.

Looking at top use cases from an industry standpoint (see diagram on next page), the picture becomes nuanced. For example, in the healthcare and technology industries, research and business intelligence rank just behind customer support. Marketing use cases show notable traction in healthcare and services.

Top agent use cases by industry

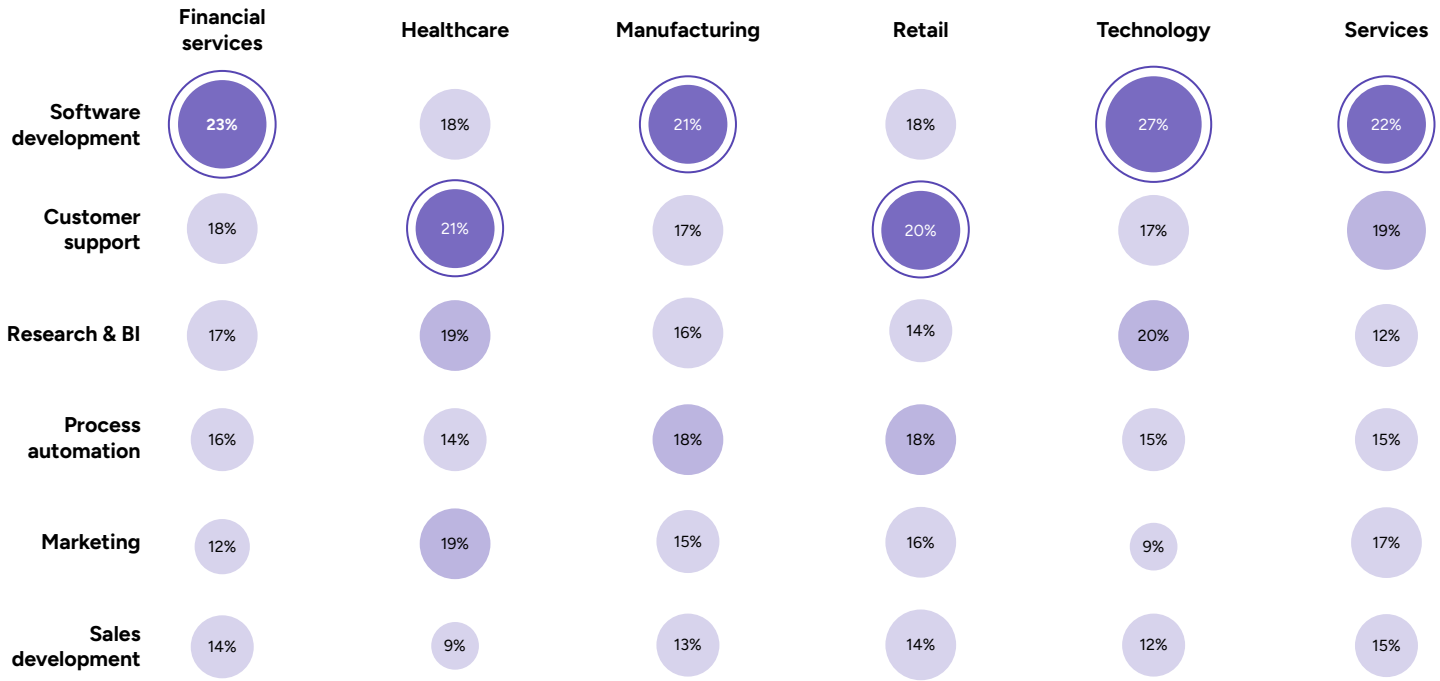

Source: G2 August 2025 Survey

However, for those using agent builder platforms, while customer support is the top use case, software development trails. This is likely because agent builders offer no-code solutions, where business users can stand up a solution with natural language — provided they clear goals and guardrails.

If you think agents won budget share in 2025, just look to next year. Over half of respondents in North America and the Asia Pacific (APAC) region expect to increase their spend in the next 12 months, while just one in three Europe, Middle East, and Africa (EMEA) buyers expect higher spend next year.

In our data, we applied an NPS-style framework, segmenting respondents into promoters, neutrals, and detractors. Certain industries, such as financial services, technology, and retail, seem to be the most bullish about future spending.

Top agent builder use cases: customer service and process automation

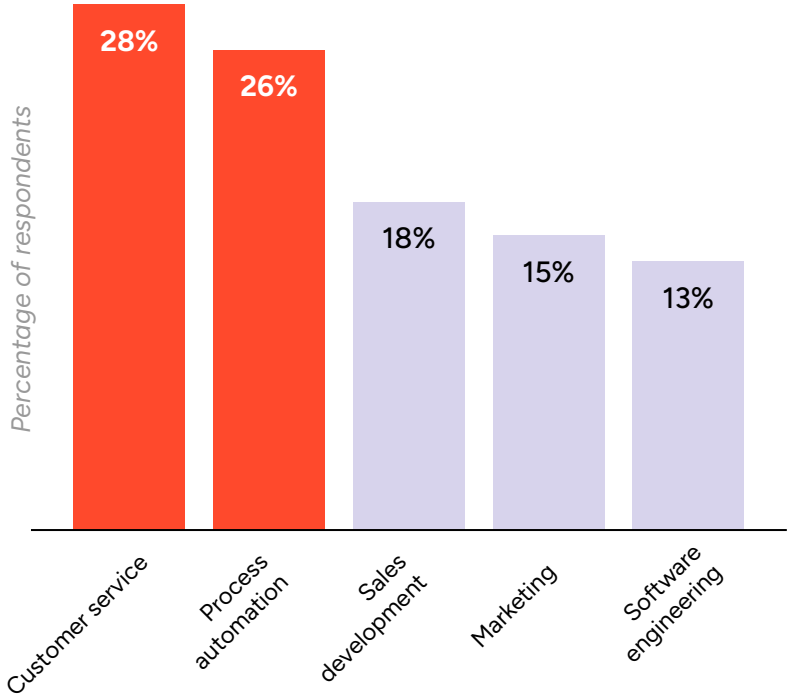

Source: G2 Review Data - August 2025

Healthcare, the industry with the biggest $5 million budgets in 2025, seemed to be taking a breather next year. For whatever reason, they saw the urgency now to get started but might be in a wait-and-see mindset before upping their level of investments.

Enterprises, many of them publicly traded, are more likely to expand agent budgets next year than their mid-market or small-business (SMB) peers. Our interviews with finance leaders at more than a dozen large companies suggest that agents enable workforce optimization, which pleases investors and reduces fixed costs in their P&Ls.

But where is all this money coming from? Are they defunding earlier AI investments?

Our research suggests that the money next year will likely come at the expense of pre-agentic Saas. Over 33% of respondents indicated that they would be willing to switch from an existing SaaS vendor (if price is equal) to acquire agent functionality.

This makes sense, as most genAI or legacy machine learning (ML) solutions lack the automation breadth that agents provide. In fact, genAI mainly provides suggestions or outputs, which require a heavy human-in-the-loop element. While that augments talent or makes them more efficient, as mentioned earlier, Parkinson's Law reminds us that work expands to fill the time allowed. That means efficiency gains alone rarely translate to enterprise-level value.

Our data suggests more heightened enthusiasm among software buyers today than late 2024, when many thought agents were at the top of the hype cycle. In our view, there is no trough of disillusionment on the horizon.

Now, this is not the result of the C-suite's FOMO. That moment has passed.

The results are rolling in, and they are more positive than most analysts predicted.

Results are rolling in

For many companies, agents are a new investment category. This development follows recent investments in genAI systems and long-standing ML programs. Already, the new bet on agents looks to be paying off.

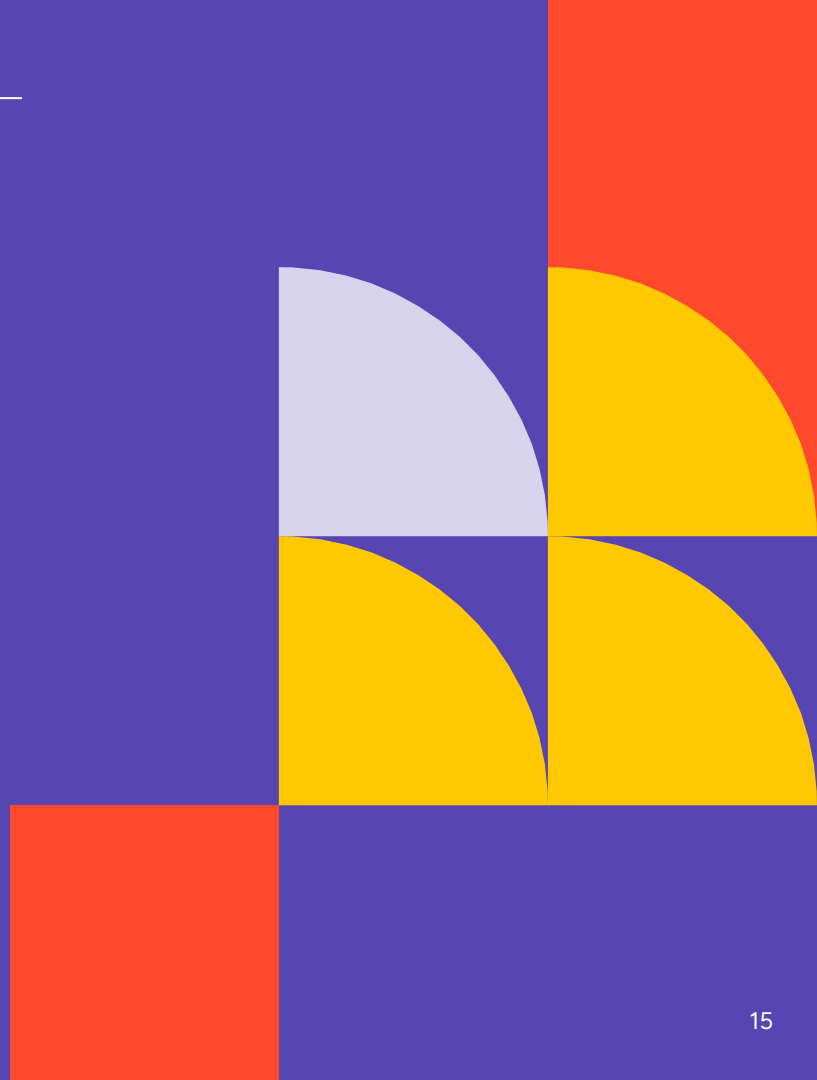

One way to frame results is “time to meaningful outcome.” Over 25% of respondents are seeing their first meaningful outcome in three months or less — truly a feat. Across all companies, the median time is six months or less. That matches or even exceeds results from prior AI programs.

When we segmented the survey data across different company sizes, SMBs get to value even faster — likely the result of fewer governance hurdles and faster deployment. They also lean on ready-made agents, which can be stood up the same day.

SMBs see the fastest time-to-value with AI agents

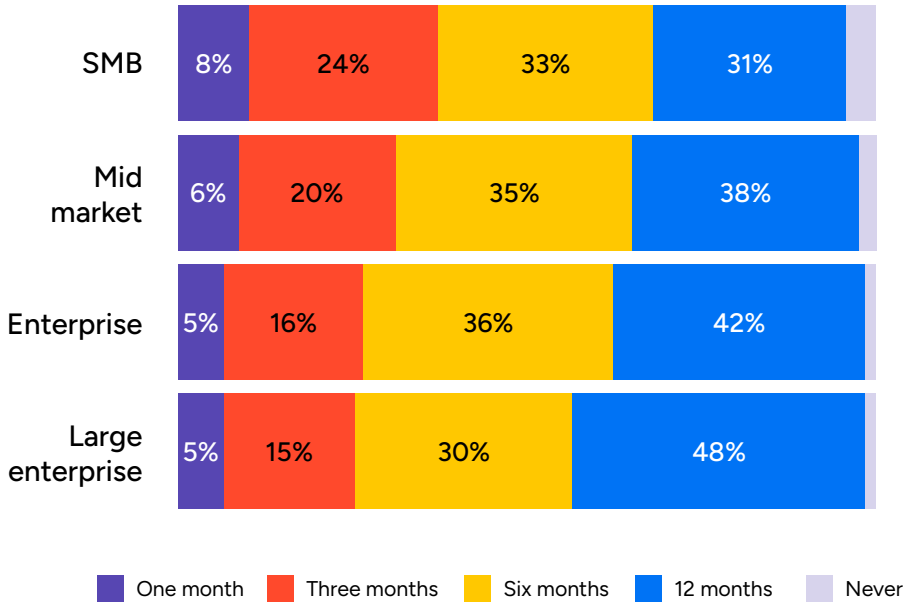

How long from agent deployment to first meaningful outcome?

Source: G2 August 2025 Survey

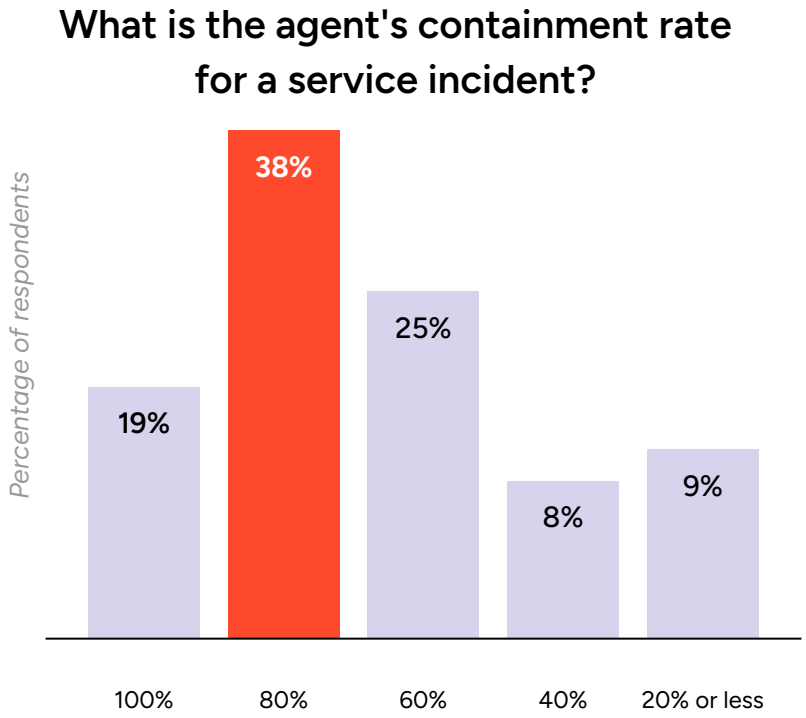

Digging into outcomes, respondents reported cost savings, productivity, and responsiveness as the top outcomes. Surprisingly, net promoter scores finished last among reported outcomes, even though customer service is the top use case.

For today's buyers, the demand is simple: Show me the money!

Respondents report a median 40% cost-per-unit savings for their most mature workflows. We focused on mature workflows, where some agents are more fine-tuned and autonomous than others. These savings exceed most productivity boost estimates for genAI or legacy ML.

Let's connect this to specific use cases G2 has tracked for over a year. Data from G2's Agent Builder category documents an 80% median containment rate for customer service incidents.

Source: G2 Agent Builder Review Data - August 2025

Now compare that to a 70% containment rate in call centers, at an average cost of $10 per incident, giving agents a clear opportunity to undercut those costs.

Note: Containment rates for genAI chatbots average ~50%, while self-service drops to ~40%. This helps explain why customer service is one of the top use cases for agent adoption.

Sales development is an emerging use case, and enterprise revenue leaders are leaning in. G2's Agent Builder category review data shows a 20% median savings in cost-per-outcome, whether that's lead generation, setting up meetings with prospects, or other top-of-funnel activities.

Given that many agentic solutions are no-money-down and pay-as-you-go, this is fast becoming a no-brainer for sales and success teams across B2B.

Agents reduce the cost per sales outcome

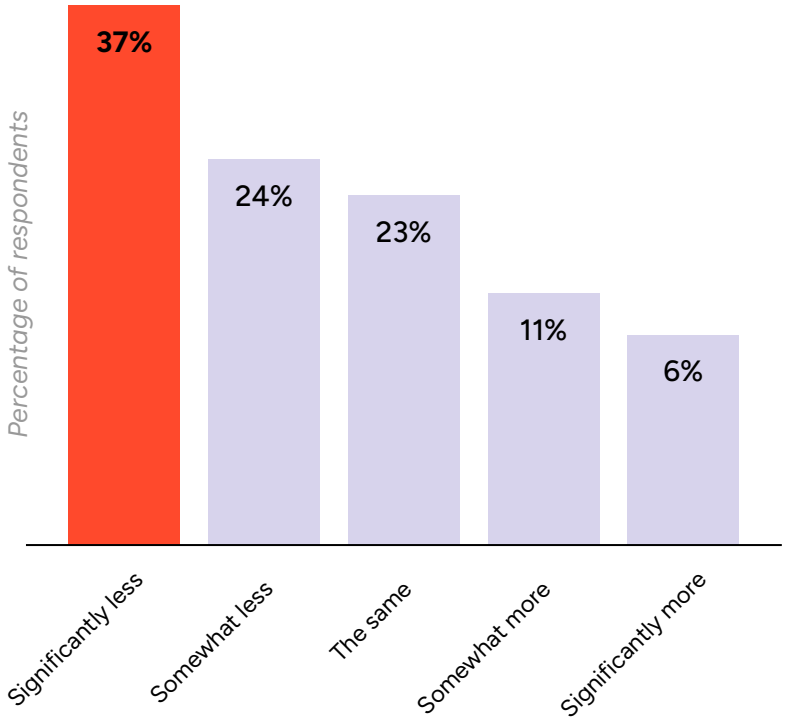

Source: G2 Agent Builder Review Data

AI agents accelerate software development velocity

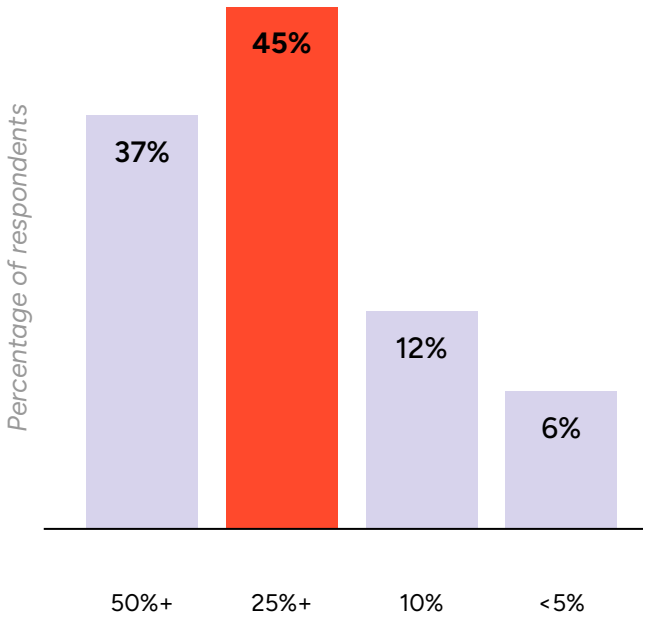

Tracking cost savings is one goal for agentic AI efforts, but the bigger opportunity lies in accelerating the delivery of products, services, and go-to-market (GTM) initiatives. That's where tracking velocity gains comes in.

Respondents to our survey reported a 23% median velocity gain for their most mature workflows. This matters because in most industries, speed to market is everything.

G2 Review Data also shows that nearly half of the reviewers reported 50% gains in marketing velocity, best measured in how quickly they reach impression goals, conversion targets, and lead-generation quotas.

For the software development use case, where retaining engineering talent is a challenge, even small velocity gains can determine whether teams meet product launch deadlines.

As you can see in this chart, over 1 in 3 companies report 50% gains, fueling faster update cycles and product launches in competitive markets.

Source: G2 Review Data - August 2025

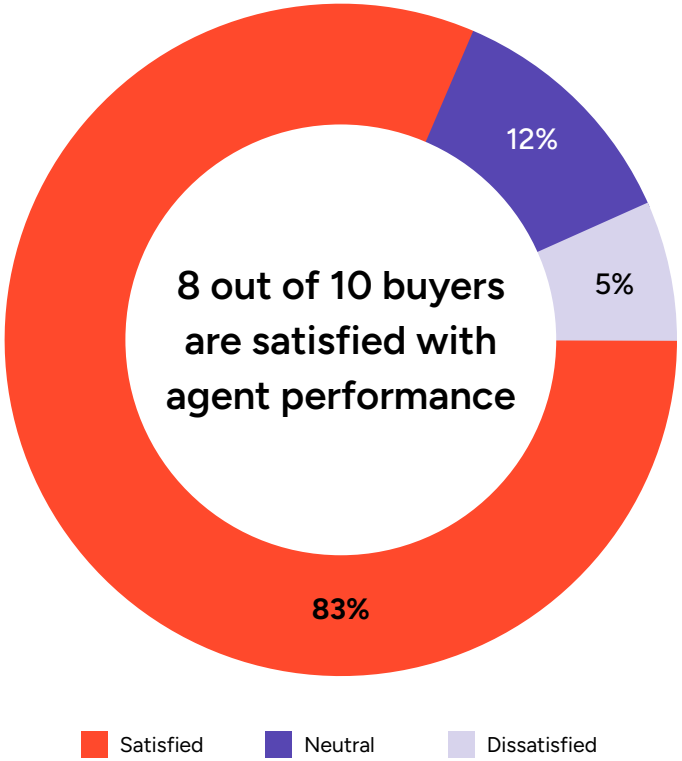

A broad use case is process automation, spanning legal, HR, procurement, and supply chain. In that instance, 8 in 10 reviewers on G2 report cycle times that are at least 25% faster.

This is arguably the biggest boost to enterprise operating leverage since the rise of robotic process automation (RPA) over a decade ago. But in this case, the agents are dynamic and handle exceptions, which RPA could not.

Taking cost savings, velocity, and other beneficial outcomes into account, it's no surprise that 8 in 10 buyers say they are satisfied with agent performance so far.

In particular, they report satisfaction with reliability, vendor support, user experience, and, most of all, value. This level of satisfaction signals a low churn or abandonment rate for agentic investments. In contrast, churn rates in many SaaS and pre-agentic AI subscriptions run in the double digits.

Source: G2 August 2025 Survey

Finally, it's not just the buyers who are satisfied. Nearly 90% of respondents report higher employee satisfaction in departments that deployed agents.

Why? Because agents take repetitive work off their plates. This tells a very different story from the headlines about job cuts and employee panic. In reality, in a world where burnout looms and enterprise employees have seen their scope of work steadily expand in recent years, agents are a welcome relief.

Together, these results offer rewards that reshape how we view agents going forward. This won't just lead to more investment and loyalty — it paves the way to higher levels of trust.

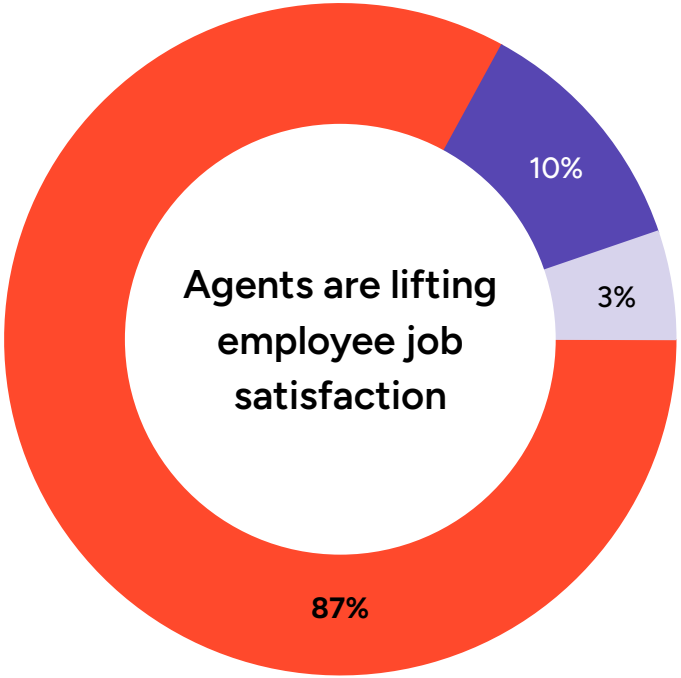

Source: G2 August 2025 Survey

Agents are gaining trust and autonomy

Despite being sold the vision of dynamic automation, nearly two-thirds of respondents report that agents need more human oversight than they anticipated. The reasons include a lack of well-structured data, steep learning curves in programming, and limited trust in the technology.

Human oversight of AI agents is still a requirement

Source: G2 August 2025 Survey

Closing trust gaps is key to adoption

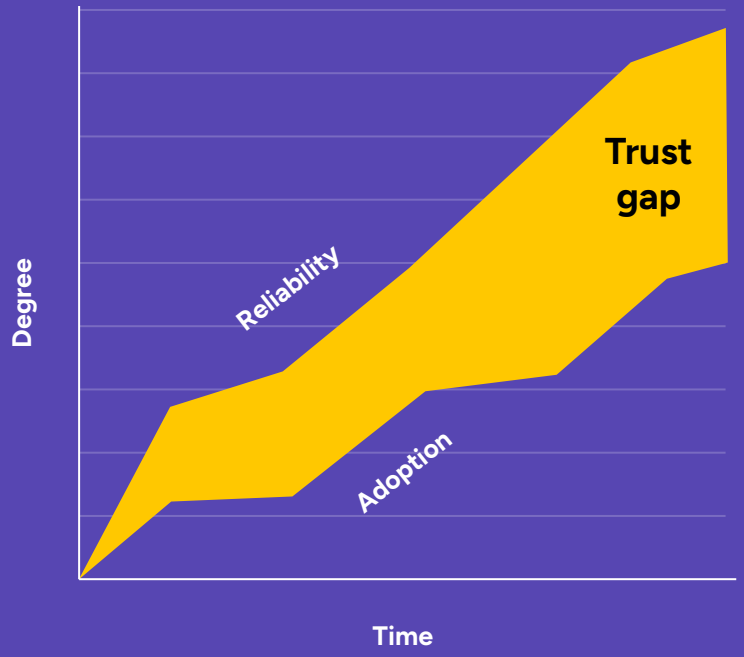

We naturally face trust gaps with new technologies. A trust gap grows when reliability improves quickly, but adoption lags because of our uncertainties or comfort with the status quo. Take cloud computing in the mid-2000s as an example. Many enterprise leaders doubted its availability, security, and governance. As a result, it took more than a decade for most companies to transition, despite the financial value and flexibility it offered.

In the case of agents, distrust is grounded in reality, not conjecture. Respondents reported security, compliance, hallucinations, and data leakage incidents over the past year. More than 25% of incidents reached major severity. Data leakage carried the greatest impact, followed by reputational damage.

These one-way doors are difficult to reverse — unlike hallucinations, which can be corrected with better data or training, or compliance issues, which can often be fixed with a system update before regulators take notice.

One might think this would halt agent deployment, keeping in mind the saying, “Trust arrives by mule but leaves in a Maserati.” In other words, it’s hard-earned and easily lost. While that may hold true in human relationships, the economic and operational promise of agents has pushed enterprise leaders to extend trust *despite* their fears.

Our interviews reveal a growing optimism, underpinning a phrase popularized by Ethan Mollick, author of “Co-Intelligence”: “Today’s agents are the worst agents we’ll ever work with. They get better every single day.” If anything, software buyers are more resolute than ever in vetting agentic solutions for performance and risk, yet still grant them as much autonomy as they can.

AI agent incidents are not trivial

Reported types of AI agent incidents in last 12 months

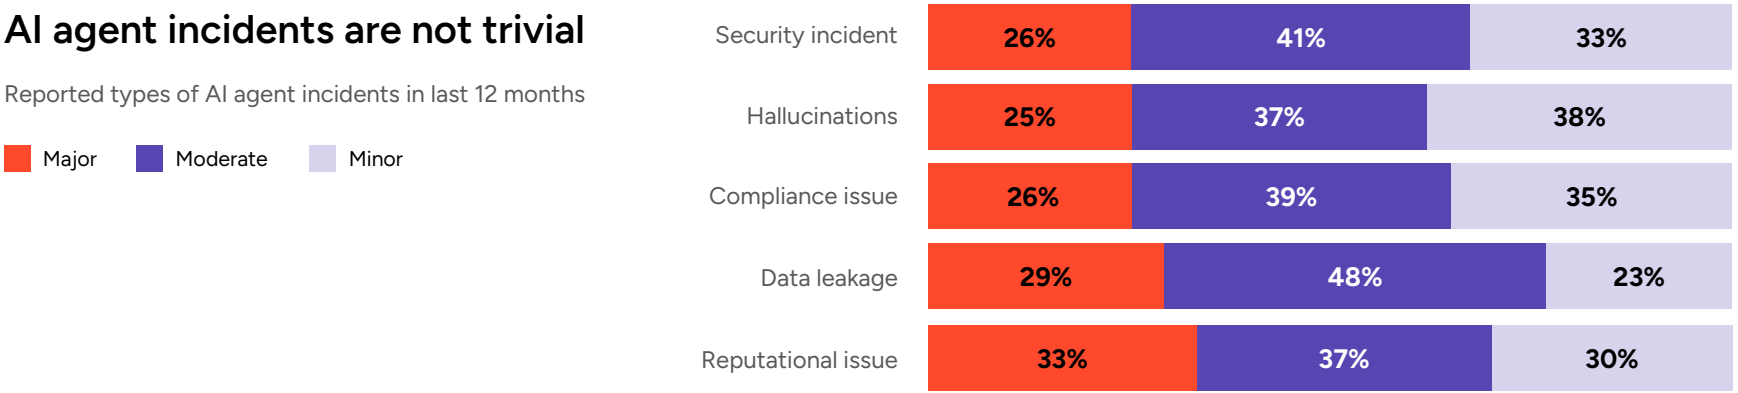

Source: G2 Survey - August 2025

Surprisingly, nearly 50% of respondents say they would grant full autonomy to agents in low-risk workflows. Yet, our research shows these workflows account for only one-third of the workload in most enterprises, while the remaining two-thirds stay human-gated for now.

So, where does that leave the state of trust as 2025 nears its end?

In the world of autonomous vehicles, trust is often discussed in terms of levels. The goal is Level 5 autonomy: no steering wheel, no pedals, full machine control. For enterprise agents, Level 5 represents the emergence of a system of agents that can create and train themselves, self-organize, and pursue human-provided goals on their own.

It's like Waymos building, programming, and deploying other Waymos — something physically impossible. But in the world of software, anything is possible.

G2 Reviews have tracked buyer sentiment toward autonomy over the past year. Nearly half of verified agent buyers place their organizations at Level 3: autonomy with guardrails. Fewer than 10% have reached a full autonomy mindset.

Total autonomy of AI agents is on the rise

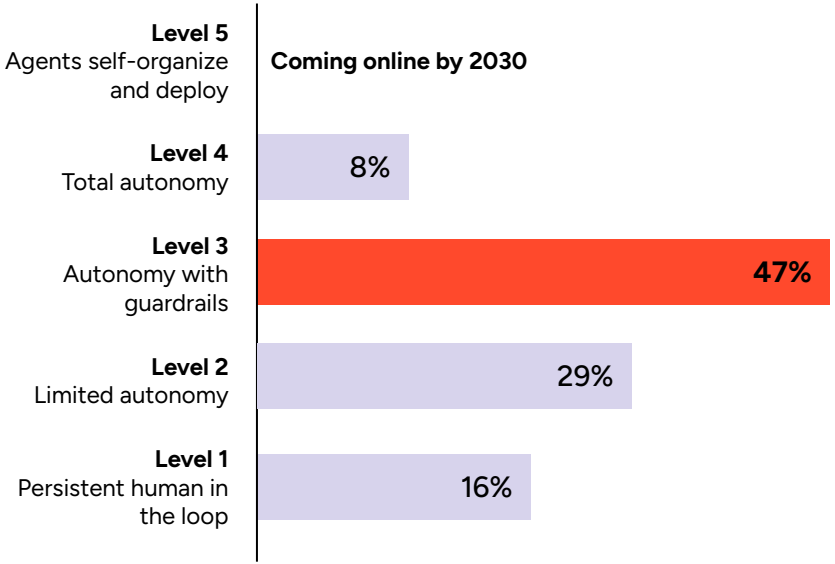

Source: G2 Review Data - August 2025

Our survey explored which agent actions users would give high autonomy to. Results show a broad range of use cases already receiving it, spanning security, IT operations, procurement, sales, and marketing.

Security and DevOps lead in high-autonomy agent actions

Maximum autonomy level allowed by organization

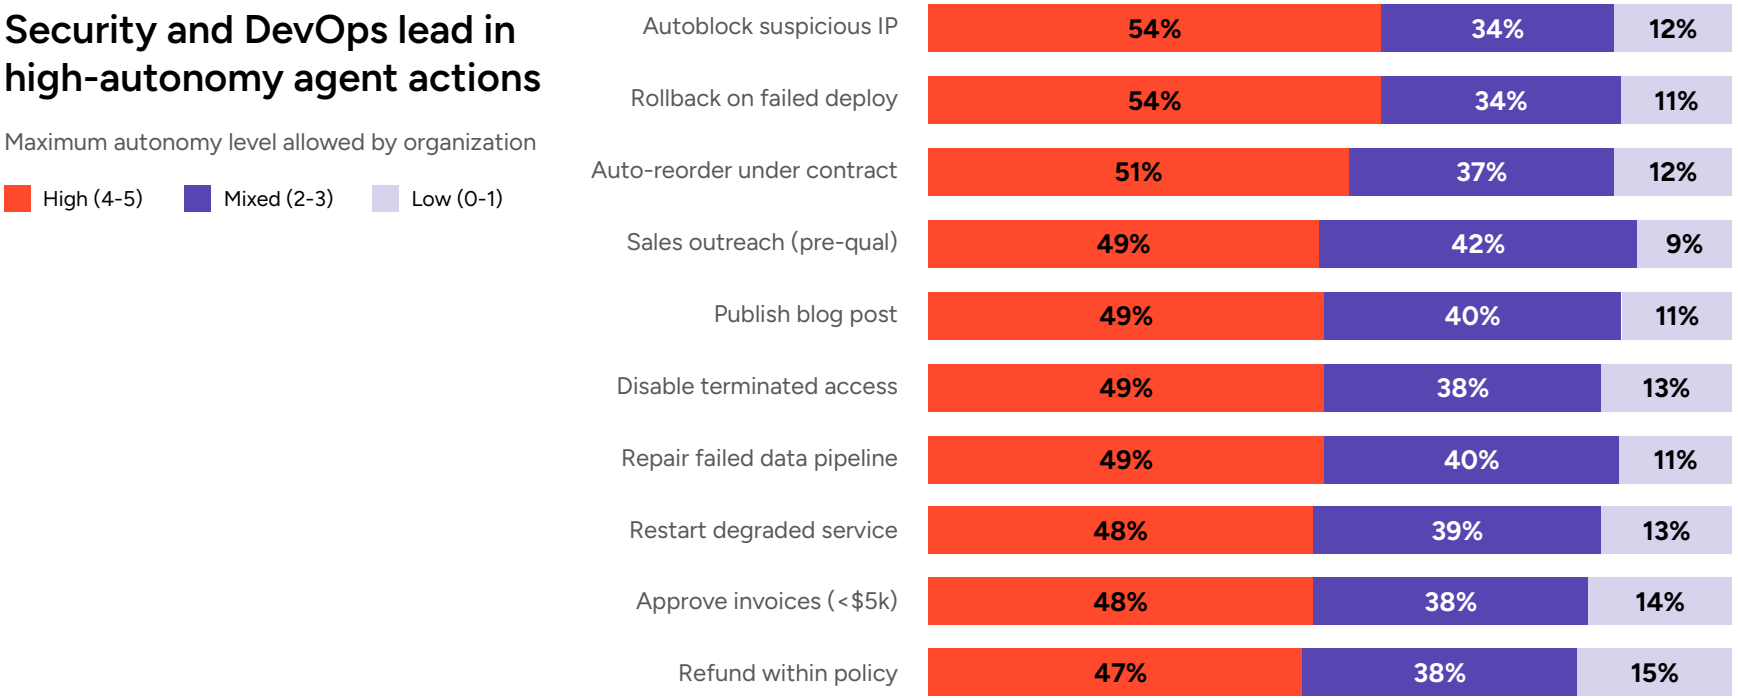

Source: G2 Survey - August 2025

Many companies have already established oversight models for agents. Remarkably, 34% follow a “let it rip” approach, where the agent takes the action, the human is informed of the result, and corrections are made after. For these organizations, pre-approval is too much of a drag on automation, making the risk tolerable.

Given the number of incidents, you might expect a pullback in autonomy next year, yet our data shows the opposite. Nearly 8 in 10 respondents plan to expand autonomy, though smaller organizations remain more conservative.

Larger companies, many of them publicly traded, face pressure to improve efficiency and bottom-line results. In our research, they consistently emerge as the most aggressive in giving agents a longer leash.

While trust is being extended to agents, there is also a growing belief that 2025 will be remembered as the “Wild Wild West” of agentic point solutions, marked by varying data quality and a lack of unified strategy. IT leaders need to rein in this unbridled enthusiasm, or at least focus it to create synergy.

That is why we believe *orchestration* will be one of the major AI trends to watch in the years ahead.

Companies of all sizes plan to increase their agents’ level of autonomy

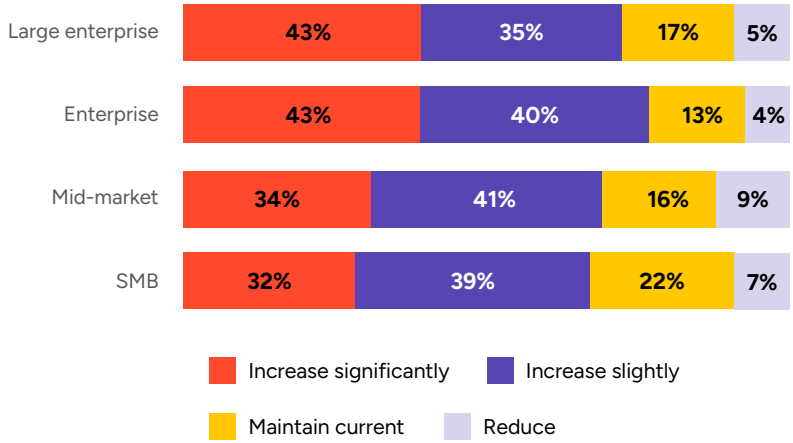

Source: G2 Review Data - August 2025

Orchestration is a major theme

Without orchestration, you don't produce a symphony. Even virtuosos without a skilled conductor can clash, creating dissonance instead of beauty.

The same holds true for the rise of AI agents. Data sources, systems, automation components, and always-on agents present an unprecedented coordination challenge for enterprises.

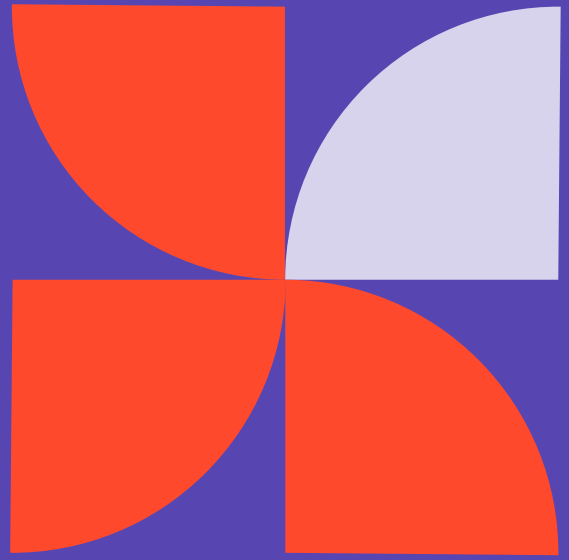

After dozens of interviews with technology leaders beginning in late 2024, we concluded that AI orchestration was an emerging theme in software. Our 2025 G2 Buyer Behavior Report confirmed the same.

It is often said that if you want to find the source of ultimate power in an organization, follow the money. So that's exactly what we did. Early genAI programs were widely reported as funded by innovation budgets and promoted by CEOs. But starting in mid-2024, recurring software budgets began absorbing AI spend.

By spring 2025, our Buyer Behavior Report survey found that 55% of AI funding came from IT departments, a share we expect to rise sharply in 2026. That's because AI has shifted from experimental projects to embedded operating technologies. Add investor pressure to test less and deliver more ROI, and the transition from R&D to IT operations makes sense.

Our 2025 Buyer Behavior Report survey also showed that IT leaders now participate in decision-making for AI purchases, partnering with department leads throughout the buying journey. This contrasts with our 2024 report, which revealed CFOs and senior leadership as the ultimate buyers.

More than half of AI solution purchases are funded by central IT budgets

Top two funding sources

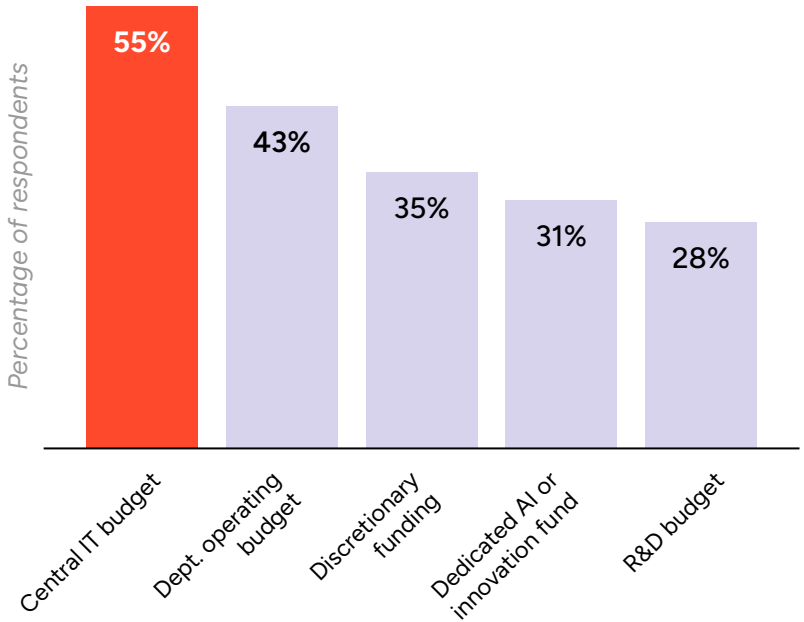

Source: G2 Buyer Behavior Report 2025

Today, IT leaders hold the baton, ensuring that every team and every part of the technology stack plays from the same page.

Data is the first concern in orchestration efforts. Coordinating and making data available across the enterprise has been a gating factor in rolling out agentic solutions. Initially, the concern was that the data wasn't clean and lacked the structure needed for LLMs and agents to produce reliable results.

While our recent survey shows a rising confidence in data hygiene, there's still a concern over the data that an agent can be trusted to act on.

9 out of 10 buyers require trusted customer data before granting agents autonomy

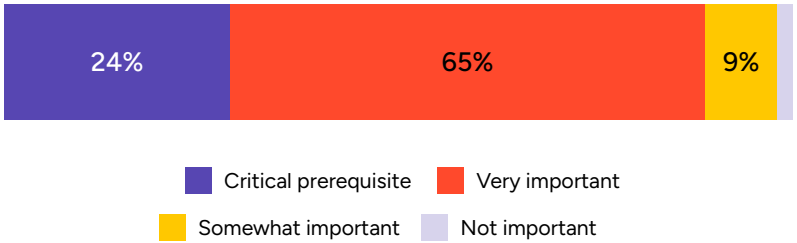

Source: G2 Survey - August 2025

Most critical functions for pre-built AI agent integrations

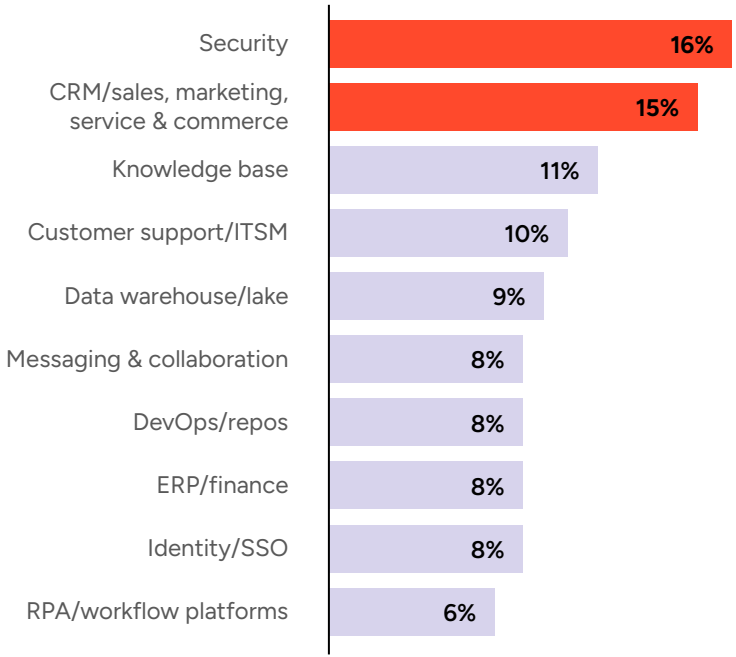

When agents interact with prospects and customers, the perceived risk is even greater. That's why 9 in 10 respondents in our August 2025 survey said having a single, trusted source of customer data was very important before granting agents autonomy. Note, this data was defined as "customer relationship management (CRM) records, support/case history, orders/billing, product-usage signals." This helps explain the success of agent builder platforms, which draw on long-standing systems of record.

Equally vital to today's orchestration efforts is an agent's ability to integrate with systems and vendor-provided ecosystems. Respondents reported that when buying agents, integration is among the top purchase criteria, along with accuracy, security, and value. Specifically, integration was a required condition when selecting agent builder platforms.

When asked about a vendor's pre-built ecosystem — including connectors and pre-configured actions and skills — the vast majority rated it a key requirement. 1 in 4 buyers said, "We wouldn't shortlist a vendor without pre-built integrations." Another 50% said, "These integrations are a tiebreaker."

Digging deeper, we asked respondents to identify the integrations that mattered most. Not surprisingly, security and CRM topped the list.

Source: G2 Survey - August 2025

Previously, integrations required engineering work or, at best, custom API development. This slowed adoption and drove up costs, even for simple integrations. Because agent success depends on access to systems and tools, complex integrations became necessary.

In November 2024, Anthropic introduced the Model Context Protocol (MCP), heralded as "the USB-C" for connecting AI solutions to enterprise systems. MCP enables instant integrations without additional programming once a server is set up. Its rapid adoption by enterprises has paved the way for plug-and-deploy agents.

Next, there's the challenge of agents communicating with each other. Initially, agents acted on systems and engaged with people, such as suppliers or customers. Technology circles have long discussed the coming agent-to-agent (A2A) era, but our research suggests it's already here.

Within the same platform (multi-agent setups), agent-to-agent communication is relatively easy to coordinate. These agents are created by the same vendor and are designed to work together. However, half of the surveyed companies report agents handing off work or exchanging messages across vendors and even across different platforms. That is a very different challenge, since not all agents speak the same language, share common goals, or operate under the same autonomy guardrails.

Agent-to-agent communications are already happening

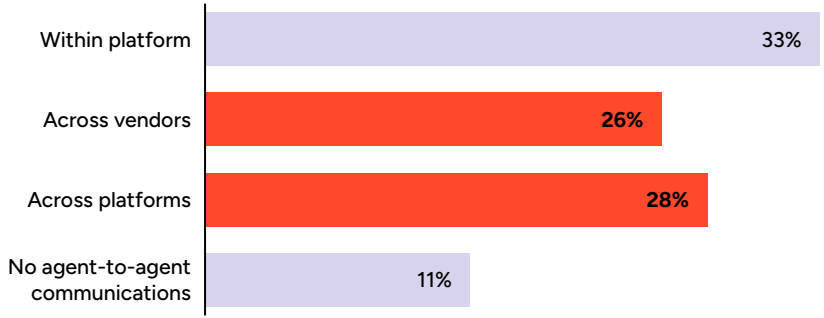

Source: G2 Survey - August 2025

The A2A protocol, like the one Google brought to market in April 2025, offers a universal language for agents to communicate, though the technology is still in its early stages. Frameworks like LangGraph and BedRock provide the stage for those conversations to happen in a coordinated way.

Finally, IT operations face the challenge of folding agents into the company's automation stack. Agents are unique because they are dynamic, probabilistic, and can reason rather than just follow rules. They adapt to edge cases and unforeseen situations. They don't grind to a halt when situations veer off script.

Earlier forms of automation, such as RPA or rules-based systems, are brittle when faced with novel scenarios. They are programmed on experiential data and cannot improvise. As a result, they are highly reliable. However, they fail to address a growing share of workflow realities and can solve only narrow situations.

While agents offer broader capabilities, IT operations leaders we interviewed cautioned that their underlying LLM technology remains prone to hallucinations. One noted, "These errors compound, especially in multi-turn or multi-step actions."

That's why IT leaders have been orchestrating a hybrid approach to automation, where RPA and rules-based systems serve as the reliable foundation of the stack. Think of them as the rhythm section for automation, with agents chiming in to complete the picture.

This is true regardless of company size (see chart on next page.)

Almost half of the respondents report taking a hybrid approach to automation, which is now the dominant strategy for integrating agents into the stack. Qualitative interviews revealed that the longer IT teams have agents in production, the more likely they are to favor a hybrid model over a fully agentic one.

Hybrid agent models lead automation stack strategy

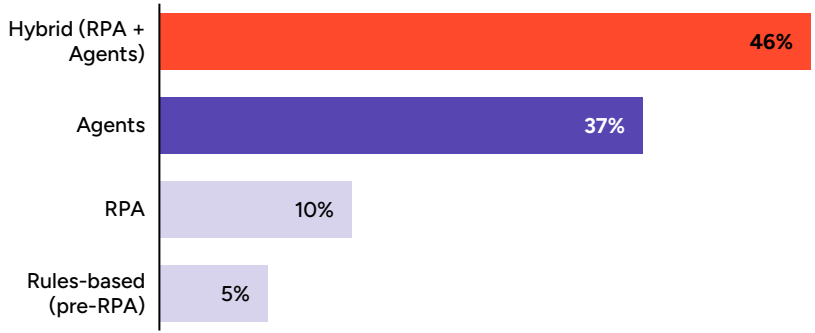

Hybrid systems pose their own coordination challenges, since mixing programming paradigms is more art than science today. At the same time, 37% of companies report a purely agentic approach to automation, which makes managing the stack even more complex. To address these challenges, vendors are offering orchestration tools designed to bring the stack together. In June 2025, G2 launched a new category to cover this emerging solution.

We expect to see a flourish of orchestration solutions emerging in 2026 to address these challenges. As IT operations become synonymous with software buyers, the pressure to coordinate before deployment will only increase.

While the choice between agent-first and hybrid presents orchestration challenges, it also highlights emerging strategies in how organizations adopt and deploy agents.

As we'll see next, a company's strategy for embracing AI agents will determine how they perform and the results they achieve.

Source: G2 Survey - August 2025

Agent strategy drives outcomes

One way to think about modernizing enterprise workflows is to compare it to upgrading a car engine. You can turbocharge it, adding forced induction and optimized airflow for higher acceleration and top speeds. Or you can hybridize it, extending the range as the electric motor takes part of the workload, producing greater efficiency and lower emissions.

Or you could choose a mixed solution like EcoBoost, designed to capture the best of both worlds. Either way, your choice determines the outcome for your car.

In the same way engine design shapes a car's performance, enterprise outcomes with AI agents come down to three choices: the automation stack that powers them, the way they're procured, and the oversight that governs their actions.

Automation stack and results

A company's automation stack strategy defines outcomes. Some IT leaders fervently believe RPA is a dependable foundation for agents and prefer a hybrid approach. But survey respondents who chose agent-first reported higher overall satisfaction with action quality. On accuracy, the agent-first approach held a 30-point lead in high satisfaction levels.

A few reasons explain this. Agents avoid the brittle handoffs of RPA bots, which often break workflows and increase failure points. Modern agents also possess self-reflection and the ability to conduct both multi-agent debates and adversarial voting. These capabilities dynamically improve quality as customer and ecosystem behaviors shift, something RPA is not designed to do.

Agent-first consistently delivered a better experience for end users and administrators alike. Not surprisingly, it also topped the list for ease of deployment. For agent-first users, orchestration was straightforward.

When it comes to security and observability, the hybrid approach shows no comparative advantage over agent-first.

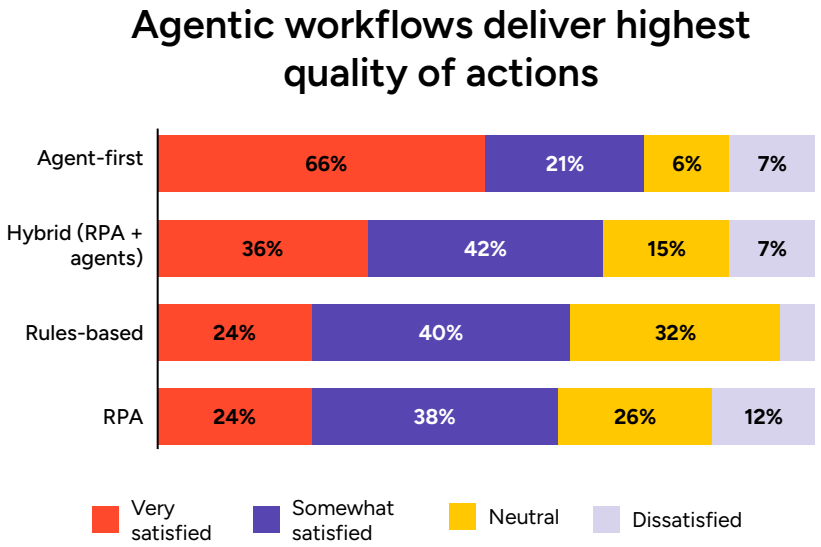

Source: G2 Survey - August 2025

Agentic strategies deliver highest satisfaction

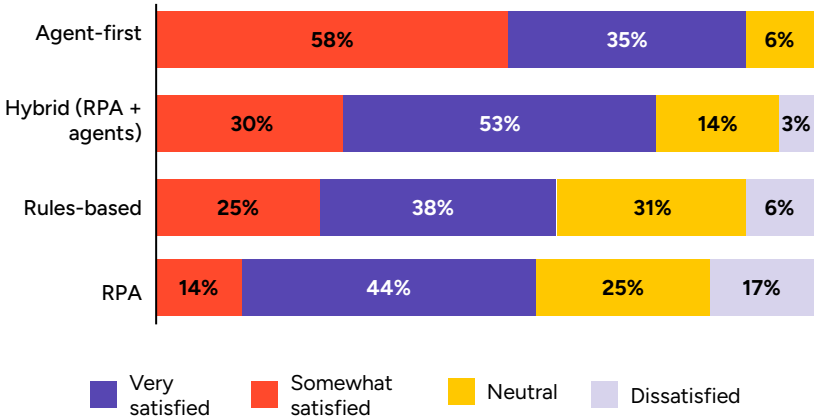

Agent-first consistently delivered a better experience for end users and administrators alike. Not surprisingly, it also topped the list for ease of deployment. For agent-first users, orchestration was straightforward.

When it comes to security and observability, the hybrid approach shows no comparative advantage over agent-first.

Looking at measurable outcomes, the agent-first approach makes high-volume workflows faster than other options, about 15 points better than hybrid on cycle time reductions. In the case of pure RPA, 12% of respondents reported either no speedup or even slower cycle times.

Finally, when it comes to cost savings, long-standing rules-based automation still leads with a median reduction of about 50%. Agent-first follows, while hybrid and RPA deliver closer to 20%. Notably, agent-first has more upside than hybrid, with nearly 1 in 5 reporting reductions above 75%. As agents mature, expect them to leapfrog rules-based automation in cost savings.

Source: G2 Survey - August 2025

It is worth noting that rules-based automation often outperforms in terms of cost savings in the most mature workflows, as reflected in our survey. These workflows have been in place for years and run at high volumes. Deterministic logic is cheaper, more stable, and requires less upkeep than UI-sensitive bots or new agent layers. This will likely change as agents gain tenure in workflows and reinforcement learning improves.

Ultimately, organizations with an agent-first automation stack reported the highest levels of overall satisfaction, and not by a small margin.

This highlights the promise of agentic technology in breaking The Iron Triangle:

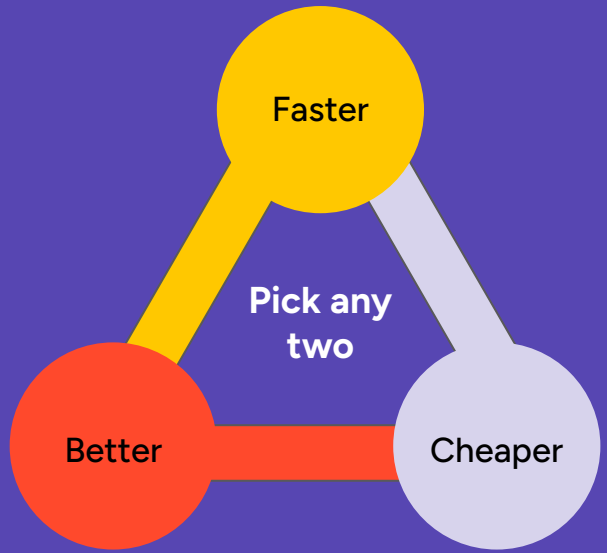

Agent procurement approach and results

The next decision is whether to acquire agents off-the-shelf, use agent builder platforms, build them in-house, or take a hybrid path with an in-house foundation and external agents.

Our survey shows 38% companies chose agent builder platforms, 30% preferred ready-made agents, while hybrid and in-house trailed.

Among the options, respondents reported that ready-made led in terms of ease of setup and deployment. Also, regarding overall satisfaction, ready-made agents received the highest top-line ratings, reflecting results across all company sizes.

Most enterprise IT interviewees leading hybrid programs viewed ready-made agents as an interim solution for department leads looking for an immediate solution to workflow bottlenecks. Their long-term intention is to build the majority of agents in-house. Taken together with survey responses from enterprise IT, the future points to a showdown between agent builder platforms and in-house builds at larger companies.

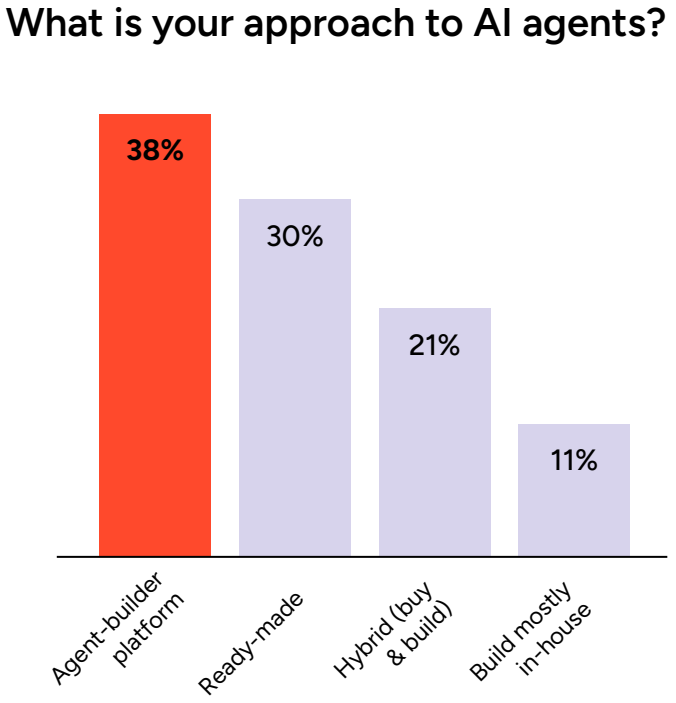

Source: G2 Survey - August 2025

The agent builder platforms deliver a trusted system of record, faster deployment, and integration with SaaS platforms. The in-house approach offers customization and greater control over data and costs over time. It is the classic buy-or-build decision.

Between the two, agent builder platforms emerge as the clear winner. Not surprisingly, respondents rated agent builder platforms much easier to deploy than in-house builds, which can become sprawling projects.

One notable finding from our survey is that 1 in 4 organizations using in-house builds have yet to see meaningful outcomes. Given their up-front costs, this delay in time-to-value is frustrating to finance leaders and senior executives.

Overall, the agent builder platform approach was nearly twice as likely to exceed expectations compared to in-house, indicating a much higher level of satisfaction. It is no surprise to see it leading the way for enterprise organizations in 2025.

AI agents are driving tangible results in as little as one quarter

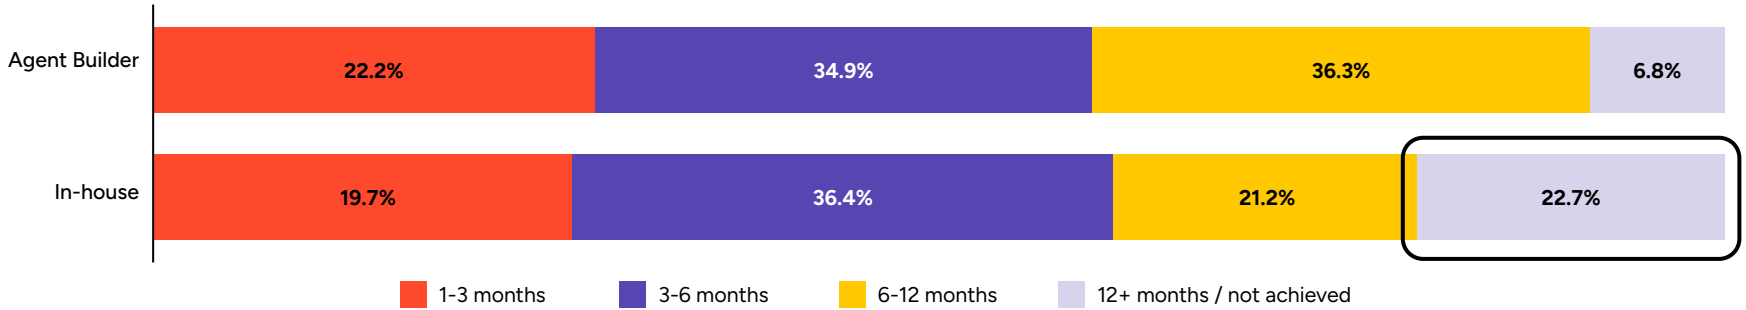

Time between kickoff and first meaningful outcome

Source: G2 Survey - August 2025

Oversight models and results

The final element of agent strategy is the level of autonomy organizations extend to agents. Our analysis focuses on two models: agent-first (human-informed) and human-gated (human-approved). By human-approved, we mean workflows with defined action guardrails or approvals required for each action.

Not surprisingly, human-gated agents delivered slightly higher accuracy. However, in terms of security and compliance, respondents reported a dead heat. It appears that the dynamic nature of agents is on par with guardrails or approvals in managing these risks.

When measuring cost savings, both approaches yielded a median of 38%. However, the human-gated approach was twice as likely to deliver savings of 75% or more.

That's because human oversight builds the trust required for organizations to deploy agents at scale, and it's at scale that the largest savings are realized.

$$\frac{\text{Efficiency}}{\text{Accuracy}} \times \text{Deployment} = \text{Agent Value}$$

Value is measured by savings or speed.

In our qualitative interviews, we were given examples of human gates, such as approving a mortgage loan as the last step or requiring sign-off prior to a digital advertising campaign going live. In both cases, an errant agent action would carry high mitigation costs. This likely played out in “let it rip” organizations that sought to maximize autonomy without considering repair costs.

The same phenomenon affected cycle time speedup, with human-gated workflows delivering 20% better net results. Survey respondents were clearly considering net results after corrections.

Ultimately, the human-gated approach delivered slightly higher agent performance satisfaction. This signals that in 2025, agents are still at a stage of development where a human in the loop is needed to achieve the best outcomes. This is likely due to the current accuracy of LLMs and the experience level of those programming and deploying agents.

Across these three strategic frameworks, organizations will likely evolve their approaches over the coming years. For many, depending on their industry or risk profile, agility will be the key to success in leveraging agents.

This attribute is important not just for those buying agents but is equally critical for the vendors selling them into this dynamic market.

Vendors must adapt to win

When cloud computing came to market nearly 20 years ago, vendors were challenged to adapt at every level. They had to shift GTM messaging from strategic transformation to hands-on technical outcomes.

Their trust signals also had to evolve from ROI projections, often delivered in slide decks, to technical proof points such as uptime, latency, and SLAs.

The entry point at buyer organizations shifted from senior executives, who held the budgets, to IT operations and technology leaders, who felt the lift-and-shift pain.

Most of all, cloud vendors had to adopt a pay-as-you-go hosting model — technology as a service. Previously, on-premises computing depended on upfront capital expenditures, perpetual licenses, and maintenance fees.

Those that adapted quickly gained a foothold in the market. Many remain industry leaders today, among the most profitable companies in the world.

When it comes to the rise of AI agents, as baseball legend Yogi Berra has said,

"It's déjà vu all over again."

Messaging must evolve

Although agents represent a once-in-a-lifetime technology breakthrough, companies selling these solutions must check hyperbole at the door.

Over 70% of respondents said the public narrative about agents is overhyped compared with the results they've seen so far. Respondents in financial services and technology felt strongest in that regard. That's why it's important to be measured in your claims about what this technology can accomplish and to be transparent about its current limitations.

Only 48% of buyers trust vendors' messaging on the capabilities and reliability of their agents. Therefore, it's important that your use cases are truly applicable to your buyers' operating environments. Focus on current reliability and accuracy metrics, not ambitious roadmaps or product investments that will "change everything next year."

Large enterprise buyers rated vendor-provided internal reliability metrics as the least trustworthy signals. To win them over, you'll need to bring third-party proof to back up your claims. That proof can come from independent review websites, analyst reports, and evaluation-benchmark providers.

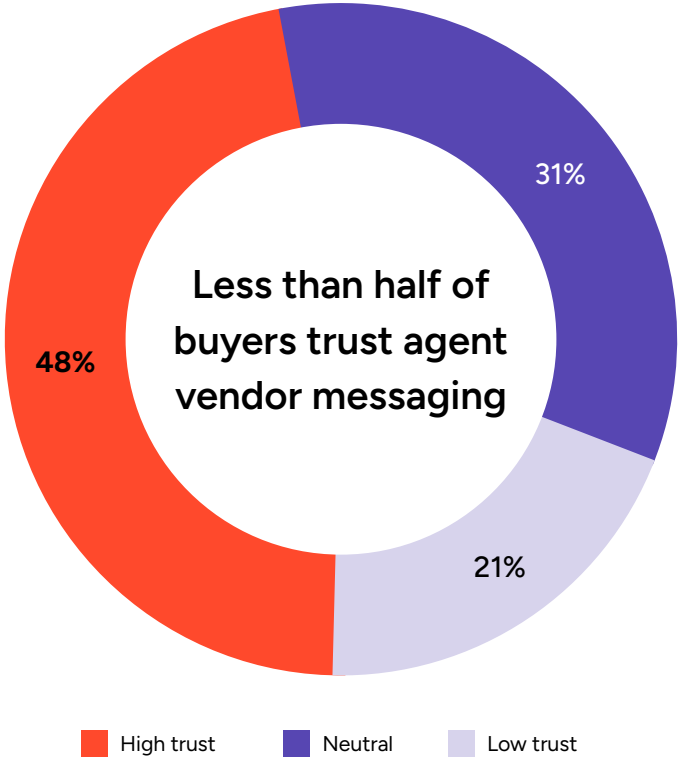

Source: G2 Survey - August 2025

Agent builders would switch software providers to gain AI agent functionality

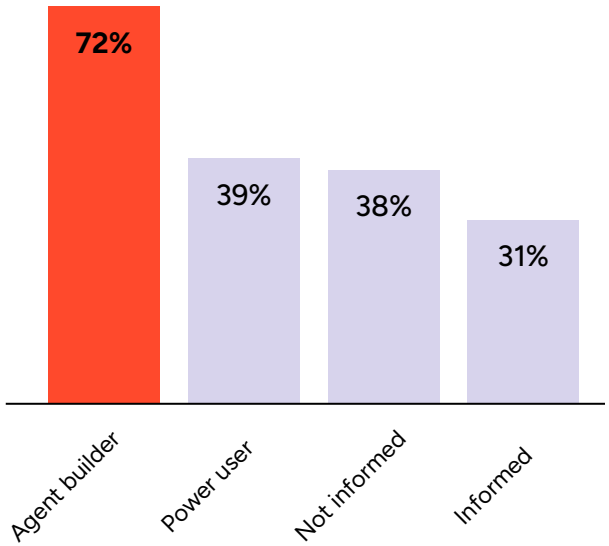

Surprisingly, the top trust signal for buyers across the board was explainability — how the technology works. If they understand your agent system step by step, they will be more comfortable granting it autonomy.

This is a big opportunity for vendors in both messaging and content development. Much like frontier labs such as OpenAI, you should publish a system card that explains how the agent is built, tested, and governed, including its known limits and safety controls.

Produce one-pagers that walk buyers through the agent process — Sense, Retrieve/Recall, Decide, and Act — while being transparent about underlying models, data, or third-party software.

While much messaging targets firmographic traits like company size, role, or industry, identifying prospects with firsthand experience of your technology can give you an edge in a competitive selling situation.

Our survey asked respondents to classify their level of experience with agents. Agent builders were almost twice as likely as other respondents to switch SaaS vendors for agents. They had a visceral experience standing up an agent, deploying it, and seeing the results firsthand.

Source: G2 Survey - August 2025

Once you identify agent builders in your target organizations, you need to speak their language. Assure they'll have live visibility into agent actions and control over them every step of the way. Give them task completion metrics and recovery measures for when things go wrong. Whenever possible, provide them with an agent health dashboard where they can monitor progress at a glance.

Agent builders have likely experienced budget overruns before, so answering the question "What's the cost per resolved task?" is critical. Give them a calculator to tune costs for variables such as model choice, context size, retries, and more.

They care most about evaluations and benchmarks. Our research shows that third-party labs or crowdsourced evaluations (e.g., LM Arena, Hugging Face, etc.) deliver stronger trust signals than customer references. This is true not only for agent builders but also for buyers in IT roles. Tool use and function-calling reliability were rated as the most important benchmarks for predicting action reliability.

Lastly, don't lose the no-code agent builder in technical jargon. They leverage agent builder platforms with chatbot-like natural language interfaces to deploy their agents.

They only need to express goals and use templates and preconfigured skills to stand up their agents. Focus on unit economics, integration breadth, and dashboards.

IT operations is the beachhead

IT operations leads all departments in agent adoption by a wide margin, at 75%. That's higher than software engineering, customer service, sales, or marketing. Don't just chase the top use cases; focus on the buying groups that are already leaning in. IT operations has a host of under-the-radar, fully automated use cases that deliver strong value, including IT remediation and data pipeline maintenance, to name a few.

Our research shows that IT operations is the department most likely to expand its budget in 2026 by a wide margin. They carry technical debt that agents can address and are adept at programming agents and understanding their performance metrics. Our data also revealed that they have the most robust datasets to support agent implementation.

IT operations is most likely to expand AI agent use in 2026

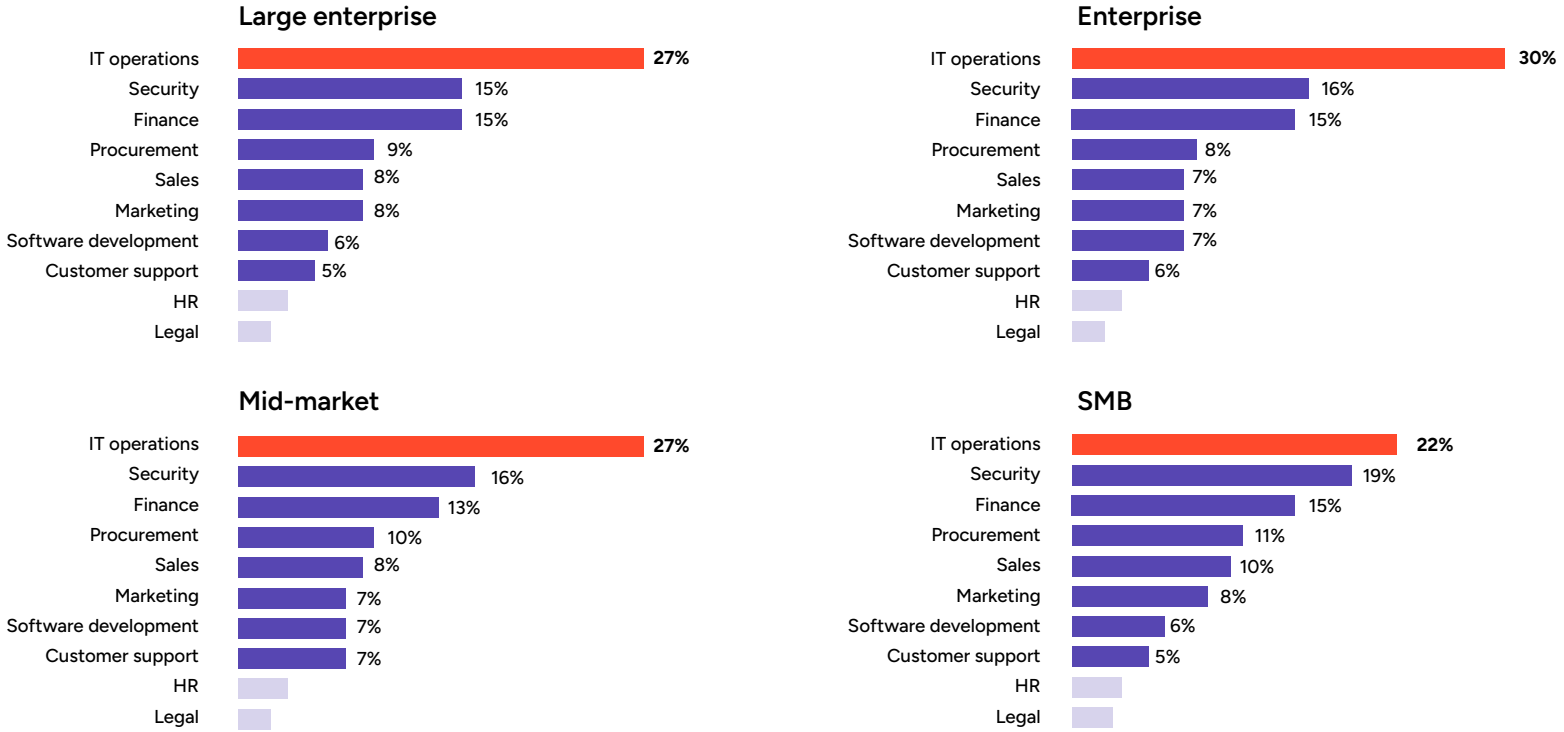

Source: G2 Survey - August 2025

According to C-suite respondents,

IT operations had the greatest potential to be transformed by fully reliable agents.

This signals ongoing support for funding requests, making IT operations a top prospect for your first outreach.

However, this will not come easily. IT operations ranks as the department with the greatest concerns. They are also the least tolerant of errors from agents. To win them over, you'll need to overcome doubts about decision accuracy, action quality, data security, and integration complexity — their top concern. They will respond more to technical signals than to satisfaction statements or case studies.

In "The Challenger Customer," researchers describe a mobilizer archetype called the "skeptic." They appear to vendors as blockers with myriad concerns, but really, they're trying to de-risk change by pressure-testing claims, forcing clarity, and pacing adoption through incremental wins. Many IT operations buyers fit this description, unlike more easygoing targets in customer service or GTM. But if you can win over these skeptics, you'll hold a winning hand in your land-and-expand efforts.

Product pricing must be dynamic

To understand why some buyers still haven't adopted agents, we explored the root causes. We asked respondents, "Have you purchased agents in the last 12 months?" If they responded no, we dug deeper to find out why.

High costs were listed as the number one reason. Expensive subscriptions, seat licenses, upgrade charges, and one-time fees all served as deterrents. Last year, vendors that approached selling agents with the "no money down, pay-as-you-go" model won more business.

Almost half of the buyers preferred variable pricing for agents over subscriptions. They saw it as the fairest way to pay for these solutions. Variable pricing includes consumption, conversations, actions, and outcomes. Our data showed that procurement favors outcome-based pricing over subscriptions, while IT prefers paying based on consumption or actions.

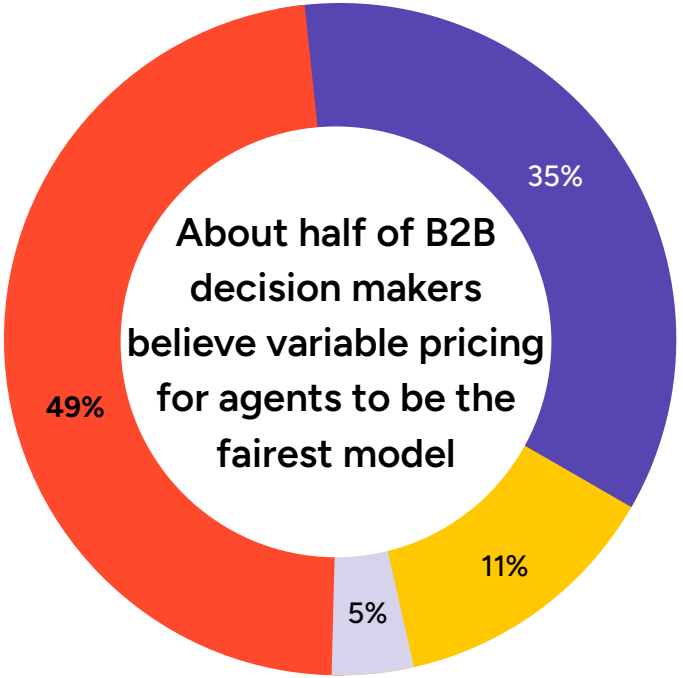

Source: G2 Survey - August 2025

A flexible pricing scheme that mirrors how cloud went to market is the key to converting non-agentic adopters and retaining customers for years to come.

When it comes to agents, SaaS-era metrics like annual recurring revenue (ARR) based on contracts will give way to technology-led growth forecasts. In this case, customer activation and continual agent improvement form the glide path to predictable revenue.

While this may seem like a lot of change management, especially in pricing models, it points to a brighter future for vendors — a future with much faster sales cycles, lower cost of sales, and no hard churn when a contract expires. Product-led growth (PLG) will not be an aspiration in this industry. It will be a reality for the most successful companies.

The future is bright for all of us, not just for companies that adapt to this dynamic situation.

Conclusion

Too often, technology breakthroughs are framed as disruptive innovations. The implication is that there will be winners and losers, regardless of how the game plays out — a zero-sum outcome. Today, some worry that agents are coming for our jobs. Others predict that, just like the game Pac-Man, agents will eat SaaS — and, for that matter, all of software as we know it.

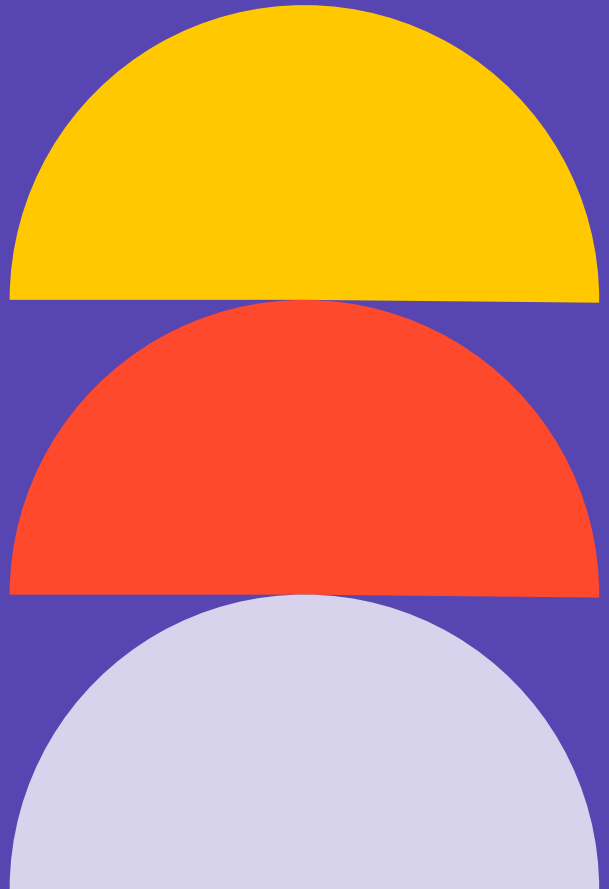

With the research and data in this report, we offer a different narrative.

When asked to predict agents' net effect on the workforce by 2028, **45% of respondents said there would be net job growth or redeployment of talent to higher-value work.** Only 20% foresaw net job reductions. Across industries, the results were more mixed: technology and financial services showed a bullish forecast, while retail and services were less optimistic.

By 2028, AI agents will create net new roles across every industry

	Net new roles	Redeployed roles	Varies	Net reductions	Too early
Technology	31%	23%	16%	19%	10%
Healthcare	30%	15%	23%	22%	10%
Retail	28%	13%	21%	29%	9%
Manufacturing	27%	14%	22%	20%	16%
Financial services	24%	27%	21%	16%	12%
Services	22%	13%	23%	25%	17%

Source: G2 Survey - August 2025

Overall, this might play out as it did when automated teller machines (ATM) rolled out in the early 1980s. Many predicted the end of the bank teller, but in a relatively short time, their roles were transformed.

ATMs automated routine cash handling, lowering the cost of running a branch. As tellers were re-trained to focus on selling financial products, their value to banks increased. Banks responded by opening 43% more branches between 1988 and 2004. This offset the decline in tellers per branch. Overall, there was a net increase in employment.

In the short term, new opportunities to design systems for agents, build them, train them, or lead digital workforces will create new categories of employment.

Over the next few years, agents will deliver efficiencies and new capabilities that unlock business growth. This would expand job opportunities in categories where only humans can deliver judgment, taste, and high-touch sales and service.

Dispensing jobs

As more ATMs were installed in the United States, the number of tellers employed did not drop

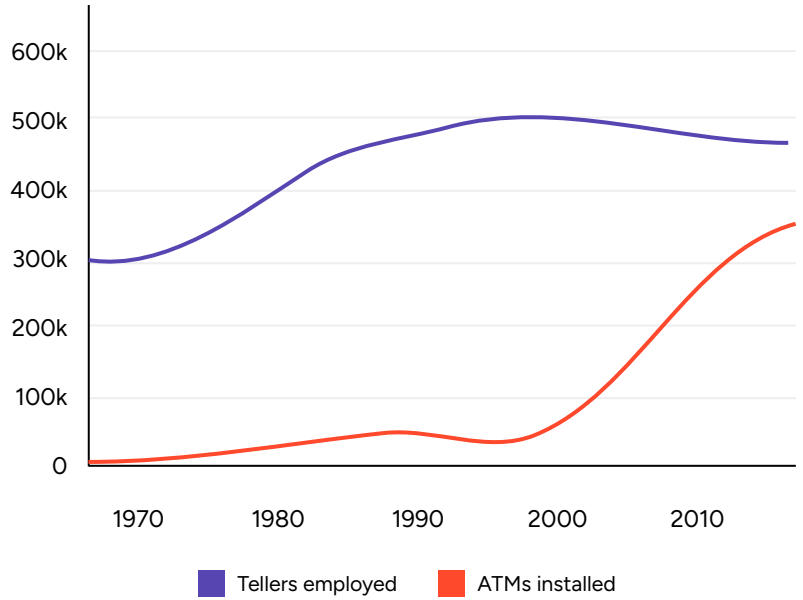

Source: IMF

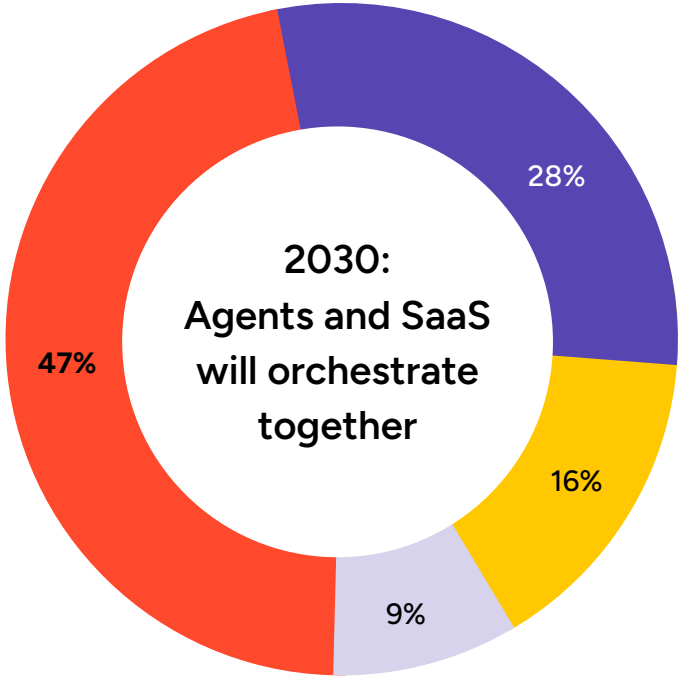

How will your organization approach software in the AI agent era?

Agents with SaaS | Agent-first | SaaS-first | Varies by function

When it comes to the future of the software industry, almost half of the respondents predict that by 2030, agents will be orchestrated over SaaS. They will coexist in a complementary relationship. That's 68% more than those who predict an agent-first world.

History again supports this thesis. Breakthrough layers often sit on the systems beneath them, as seen with mobile wallets on card rails. In other cases, they work alongside existing solutions, such as smartphones still running in tandem with PCs.

Source: G2 Survey - August 2025

At G2, we are preparing for a future in which software and human talent will be as relevant, if not more so, than ever before, and where agents help both reach their full potential.

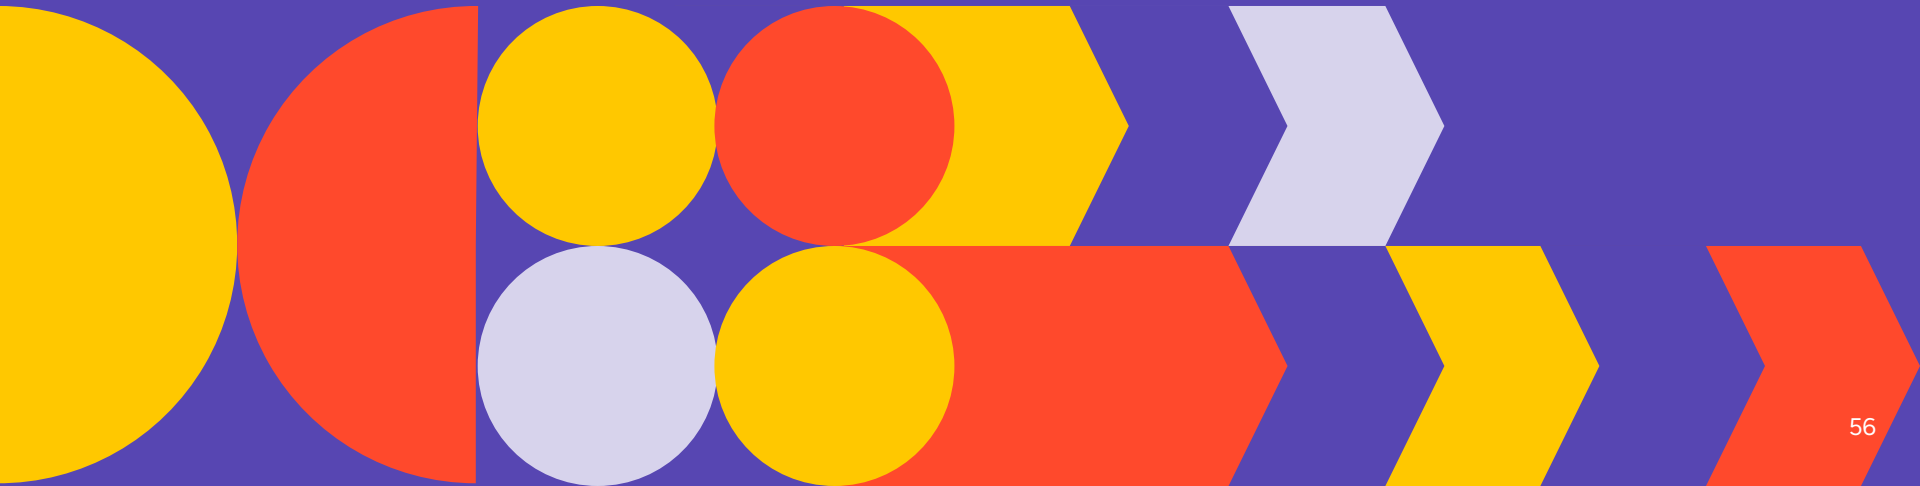

Methodology

G2 fielded an online survey among 1,035 B2B decision makers responsible for, or influencing, purchase decisions for departments, multiple departments, operating units, or entire businesses. Respondents had job titles ranging from individual contributor to manager, director, vice president, or higher.

To maximize differentiation for vendors, this survey defines small-medium business (SMB) as a company with 1-250 employees, mid-market as a company with 250-1,000 employees, enterprise as a company with 1,000-5,000 employees, and large enterprise as a company with 5,000+ employees. The survey was conducted in August 2025 and includes a global pool of respondents across North America, EMEA, and APAC.

G2 Review Data was employed to validate various findings as well as to generate insights for this report. Qualitative research was conducted through over 40 interviews with technology, IT, and executive leaders. Generative AI reasoning models were used to define the study's focus areas, optimize survey design, and analyze results to inspire writing and data visualizations.

About G2

G2 is the world's largest and most trusted software marketplace. More than 100 million people annually — including employees at all Fortune 500 companies — use G2 to make smarter software decisions based on authentic peer reviews. Thousands of software and services companies of all sizes partner with G2 to build their reputation and grow their business, including Salesforce, HubSpot, Zoom, and Adobe. To learn more about where you go for software, visit www.g2.com and follow us on LinkedIn.

Author and Lead Researcher

Tim Sanders

Creative

Daniella Alscher

Kyara Lomer-Camarena

Mike Puglielli

Contributors

Sarah Reks

Emily Greathouse

Jeffrey Lin

Alex David

Supanna Das

Jenny Gardynski

Michael Wood

Mara Calvello

Blue Bowen