Engineering Business Outcomes:

The Shift to Agentic Systems

As agentic AI scales, closing the gap between governance and operational readiness has become a top priority.



outsystems | KPMG | CIODIVE

Custom content for OutSystems from Studio by Informa TechTarget

Foreword

Over the past year, building software and building AI systems became the same job, and investment in agentic systems exploded. We've seen rogue agents hacking private networks, unpredictable inference costs, and apocalyptic predictions dominating the news cycle. But the reality is entirely different from all the noise. Companies are rebuilding critical workflows around what agents can do. And they're doing it without jeopardizing their business.

In conversations with customers around the world, I see what that looks like in practice. A manufacturing team has agents interpreting technical drawings. A travel company designs personalized itineraries for every traveler. A bank catches technical issues before customers ever notice them. An insurer checks compliance as work happens instead of waiting for a quarterly review. Inside OutSystems, agents help our own teams with research, knowledge management, and sales strategy.

This is work our customers and our employees depend on, and that shift raises questions that are different from the days of pilots and experimentation. So, we worked with KPMG and CIO Dive's Studio to survey 858 executives about how they run agentic AI today.

What they told us is worth your time. Agentic AI has moved into risk and compliance, customer service, IT operations, and finance. Confidence is high, with a vast majority saying governance is keeping pace with deployment. Yet only a little over a quarter of them feel they are well prepared for when agents come into conflict with one another. Most are managing more than 20 AI tools, models, and platforms at once. The gap between speed and control is the defining engineering problem of this era.

The ones closing that gap are doing it with architecture rather than caution. They are the pragmatists building AI into the applications and workflows that run their business. They put AI where it does real work and measure it against a clear result.

Read on to see how your peers view agentic systems, where they are finding measurable value, and what they plan to build next.



Tiago Azevedo,
CIO, OutSystems



Executive Summary

Agentic AI has moved beyond isolated pilots and into consequential business workflows, introducing both risks and opportunities. **While 82% of executives said they believed their governance could keep up with adoption, there was clearly a gap between adoption and widespread operationalization: Only 26% of leaders described their organizations as "very well prepared" to resolve agent conflicts.** Moreover, **71% expected agentic applications to change over the next five years.**

Today's challenge is governing the ever-growing number of interconnected agents in a way that's safe, sustainable and directly tied to business outcomes.

So how are leaders rising to this new operational challenge? OutSystems and KPMG worked with CIO Dive's Studio to answer this question by surveying 858 executives on their use of agentic AI and the challenges they face.

The findings showed that scaling agentic AI safely required enterprises to translate governance policies into operational controls that work consistently across agents, models, applications and data while preserving enterprise context, technology choice and adaptability.

82%

of executives said they believed their governance could keep up with adoption

26%

of leaders described their organizations as "very well prepared" to resolve agent conflicts

71%

expected agentic applications to change over the next five years.



Introduction:
The agentic enterprise takes shape

Agentic AI is now embedded in a broad range of mission-critical business operations: **73% of organizations used it in risk, compliance and security operations, 69% in customer service, 63% in IT operations, and 63% in finance and accounting.**

According to Gonçalo Borrêga, senior vice president of product management at OutSystems, that shift changes the risk profile for CIOs. Agents no longer operate as isolated assistants but as systems that can act with increasing autonomy. Moreover, they must now interact with other systems while preserving identity and permissions, sharing context, and handing work to other agents, often across different platforms. "The higher the goal is, the more coordination effort is involved, just like it is with the team. An error in one system can lead to 10 errors down the line," he said.

Michael Harper, managing director at KPMG, gave an example of how this interconnected future raised the stakes for orchestration. "There is also the reality of cumulative risk. A 98% reliable agent invoked five times in a sequence is actually much less reliable than most executives assume," he said.

At the same time, economic pressure is making AI strategy more deliberate. Borrêga linked cost discipline directly to model choice and use-case selection. "Small models with proper context management are usually better by far than a large frontier model with bad context," he said. He also warned that organizations needed a framework to understand where to use agentic AI, as well as each model. "It's not just about choosing the model. It's really about choosing the model in the context of your use case and what tools the agent can access," Borrêga said.

71%

of respondents expected enterprise applications to become more integrated, with networks of agents coordinating tasks, decisions and handoffs across multiple applications and workflows.

Agentic priorities differ by region and industry

45%

of North American organizations expected rising AI costs to drive centralized AI purchasing and strategy, compared with 31% in Europe and 28% in Asia and Australia.

40%

of European respondents expected greater reliance on strategic partners or managed services, compared with 31% in North America and 26% in Asia and Australia.

14%

of Asian and Australian organizations reported having automated controls, audit trails and human review points for agentic AI systems, compared with just 3% in Europe and 2% in North America.

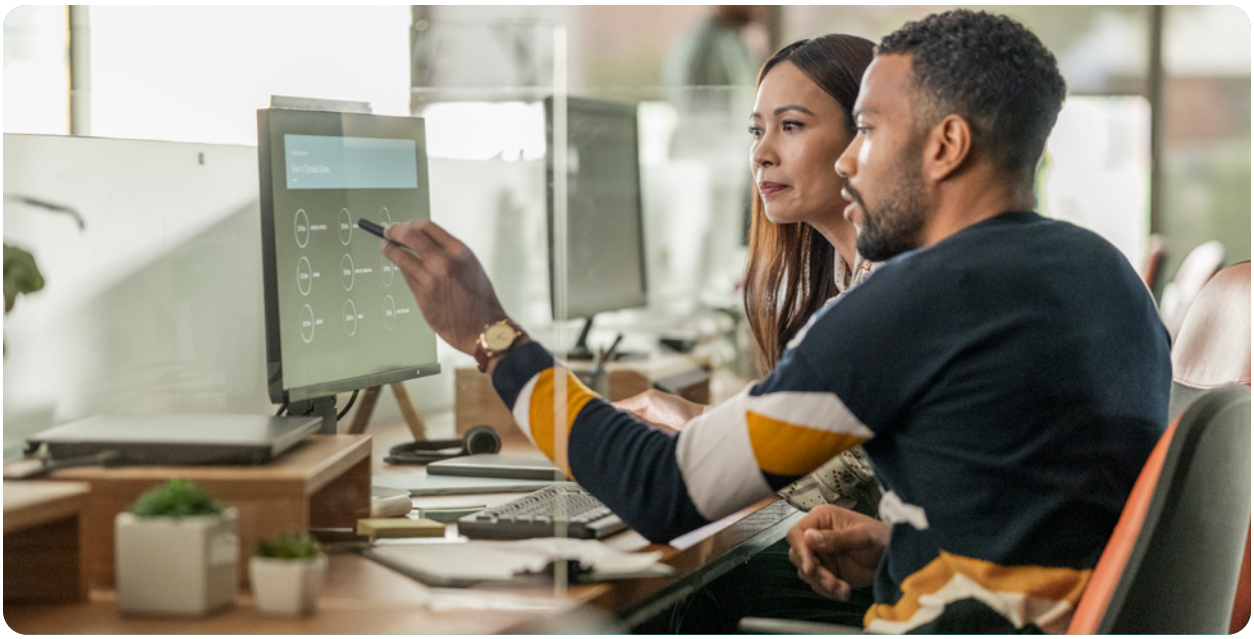
67%

of insurance respondents were very or extremely concerned about sensitive-data exposure, compared with 59% in banking and 53% in manufacturing.

The gap between perceived control and operational readiness

Executives were broadly confident in their ability to govern and control agentic AI as it scaled across the enterprise. While **82% disagreed that deployment was moving faster than governance, integration and monitoring could keep up with,** the results also found a more nuanced picture of operational readiness: Only **26% considered themselves "very well prepared" to resolve agent conflicts.** Agent collision has therefore become an operational stress test.

Control is also a critical enabler of strategy. This includes data sovereignty, cloud costs, model selection, coding agent tooling and deployment options. To be truly effective, Borrêga said, control required full-stack choice. A vendor-neutral platform, for instance, can give organizations greater control over those variables, without locking them into a single vendor.



Borrêga described the gap between confidence and operational readiness as a system-design problem. "We see in a lot of enterprises that they're building agents in silos, so managing conflicts that happen over real transactional data is a hidden problem that I would say many enterprises aren't even aware of. And it's a very hard problem to solve."

There is some legitimate basis for confidence, even though oversight remains largely procedural and human-led. For instance, **83% used existing IT, security or compliance processes; have defined review processes for higher-risk use cases; or centralized policies and monitoring.** Borrêga explained why this oversight still mattered: "Reversibility of an action is the key criterion when thinking about whether you should design escalation as a first-class path versus designing it as exception management."

Harper pointed to change control as another gap between procedural governance and true operational control in environments where modest changes could alter agent behavior without passing through the review mechanisms used in conventional software. "Tweaking a prompt or updating an underlying foundation model can completely alter an agent's behavior without ever triggering a formal SDLC gate," he said.

On the other hand, only **5% of enterprises had automated controls, audit trails and human review points for agentic systems,** which suggested that many didn't continuously govern agents once they started taking actions across systems. "You need logging for post-incident analysis, but you also need to ensure strong authentication and identity for the agent throughout the entire workflow," Borrêga said.

Harper argued that true observability must let teams reconstruct why an agent made a decision, what context it retrieved and which tools it called. "Traditional logging is built for debugging code, not for auditing AI decisions. Without that level of observability, governance is just guesswork."

In this respect, the biggest challenge is translating policy into repeatable runtime control. Enterprises still need a dependable way to apply rules, permissions, monitoring, escalation and remediation across the places where agents operate, rather than only when new systems are approved. That's why Borrêga has proposed an agentic systems platform to provide a unified control plane and end-to-end governance.

Engineering security, context, control and reliability into agentic systems

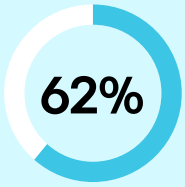
The leading priorities for enterprises are operational governance and validation. Eighty-seven percent of **respondents rated the ability to consistently apply security, compliance and governance policies as a top priority,** while **81% said the same about testing, evaluating and validating agent behavior.**

Closing the gap requires that governance become integral to how agentic systems operate, rather than something applied only when agents are approved for deployment. As Harper said, "Assurance for agentic systems has to be engineered from day one. You cannot just inspect it after the fact." He recommended incorporating controls at the point of action, policy-as-code, runtime guardrails and intentionally designed human oversight.

Ensuring agents are **grounded in a governed enterprise context was a top priority for 73% of respondents,** while **35% cited data consistency and context sharing as a major orchestration hurdle.** Moreover, **74% of respondents rated integration with existing apps and systems as very or extremely important,** while **38% considered integration with existing apps, data, and workflows** among the most important scaling capabilities. After all, agents can't reliably complete complex business work if they can only access generic model knowledge and not real business data.

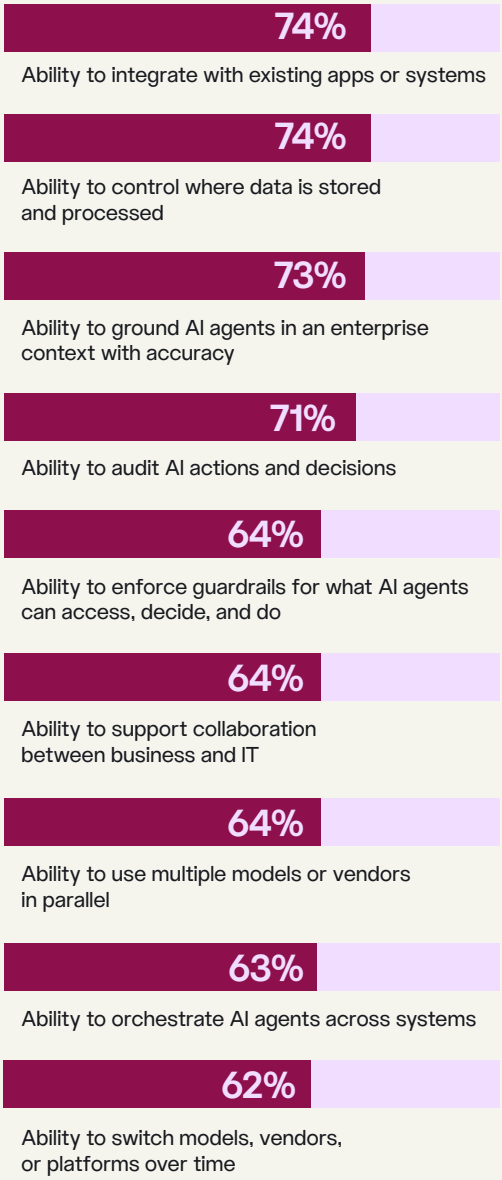


point to agent orchestration across systems,



value the ability to switch models, vendors, and platforms over time.

The capabilities that are extremely or very important when selecting platforms or tools for agentic AI are:



Borrêga said he considered shared enterprise context a prerequisite for consistent decisions. "You need to have a data strategy and a context-management strategy that ensures data is semantically ready for agents to work on top of," he said. The objective is to provide enough operational context for agents to act, while keeping access tightly bounded.

Security is one of the clearest starting points: **67% of respondents considered security, permissions and policy enforcement across agents a leading orchestration hurdle.** Borrêga said that runtime security must be deterministic even when the LLM was not. In other words, every agent execution and downstream system and API call should inherit the authenticated identity and permissions of the user or process to prevent risks such as hallucinations or prompt manipulation from expanding access.

"If you don't have that operational control over security in your agentic systems, you're at greater risk of data or policy breaches," he warned.

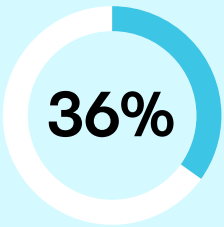
Finally, enterprises must engineer reliability across chained workflows, which **38% cited as a top orchestration hurdle.** Harper warned that individual steps could appear reliable while the overall workflow still produced the wrong result: "If an agent chains actions across disparate systems without strict data lineage, it can produce outputs that are technically correct at each individual step but materially wrong by the end of the process," he said. Borrêga therefore recommended treating every individual agent as a unit of work and then testing the complete orchestration in much the same way as software teams test individual services and their integrations.



Governing a diverse AI ecosystem without sacrificing choice

Engineering the agentic enterprise is a challenge when organizations assemble increasingly diverse AI ecosystems. Seventy-two percent **of organizations had more than 20 distinct AI tools, models or platforms in operation.** The highest degree of sprawl appeared in the IT, technology and insurance sectors. However, even at the lower end of the range, a piecemeal approach to managing agents still introduces a challenging operational overhead. When that happens, it only becomes harder to manage governance, costs and interoperability.

Excessive tool sprawl, however, doesn't mean leaders should attempt to force-fit a wide range of point tools into a single, restrictive suite. Rather, the goal should be to unify engineering, orchestration and governance to preserve flexibility across the broader environment. Such a strategy could benefit the **36% of organizations surveyed that said they had a fragmented approach to deciding where and how to use agentic AI,** which is due in part to lacking that unified strategy.



of organizations surveyed have a fragmented approach to deciding where and how to use agentic AI

According to Harper, central IT and security teams should own policy, risk appetite, shared platform capabilities, observability standards, and approved model and tool catalogs, while business units own use-case prioritization, domain-specific design, adoption and process redesign. "The sweet spot is a thin, opinionated central platform that makes the safest path the easiest path, giving business units the autonomy to deliver value securely," he said.

Borrêga framed this unification as "agentic systems engineering," which he said was, first and foremost, an architectural discipline. "Even with a platform, how you design and architect the system really needs to take into account logging, security, traceability, identity, context and evaluation at the unit level. Then you orchestrate on top of that with humans in the loop," he said.

This is another area where full stack comes into play. Organizations must have the flexibility they need to adapt continuously to the AI landscape, given its constant state of flux. After all, the best tool or model for a given use case today might not be the best option tomorrow. "It's unlikely, when you move from one agent or one task to multiple agents and multiple tasks, that all those agents are going to be in the same platform," Borrêga said. Agentic systems engineering includes cloud and on-premises or regional deployment and architecture options where sovereignty, regulatory or industry requirements demand them.

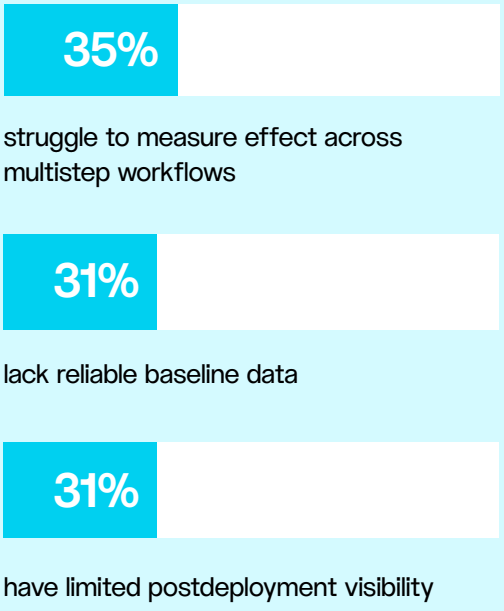
Preserving that flexibility also depends on being able to change agentic systems safely. Harper suggested that strong operational controls could make teams more willing to experiment. "Teams that know they can reliably roll back an agent to a previously known good state in a matter of minutes are going to experiment much more boldly than teams terrified of breaking production," he said.

That is, for the most part, what control is all about, and it's what gives organizations the choice they need to ensure appropriate levels of security, observability and reliability for any given workload. That unbiased choice of LLMs also keeps costs more predictable and prevents vendor lock-in, while giving teams the tools they need to innovate faster.

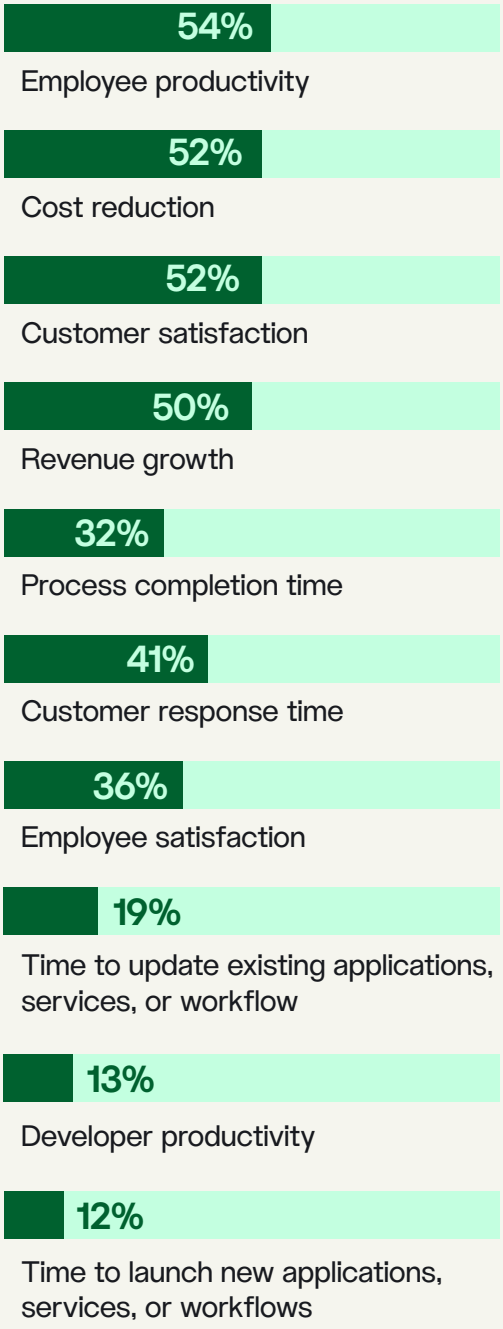
Measuring business value across agentic workflows

As agentic AI moves deeper into business operations, the measures used to judge its success are also shifting beyond productivity gains and cost-cutting. While **54% and 52% of respondents cited employee productivity and cost-cutting gains as key metrics for evaluating the effect of agentic AI, 64% cited error reduction and 60% cited compliance and risk reduction.** Other measures of success included r**evenue growth (50%) and process completion time (43%),** while just **13% cited developer productivity** as a key success metric.

The top barriers to measuring the value of agentic AI



Metrics companies use most to evaluate the impact of agentic AI on business outcomes:





Borrêga argued that leaders should judge the value of AI by business outcomes. "You could pick a decision-making process that's high value for the company rather than an easy-to-implement productivity use case that improves customer service, but in the end, that might have you paying for something that people are still going to spend a lot of time reviewing afterward. If so, it might not add much value, if any."

That's why outcome measurement must span the entire workflow. For instance, an isolated task might appear technically successful in a pilot environment, but the overall end-to-end workflow could still end up being costlier, slower, or riskier than it was before. "Boards and CIOs should expand their view of how AI is performing by measuring and reviewing process-level KPIs as well as token consumption and costs," Harper said.

Borrêga recommended establishing a baseline for existing processes and measuring the change after redesigning that process with the use of agents. "How do we benchmark the new system versus the old way of doing things?" he asked. "The individual use of AI, by contrast, is very hard to relate to the outcome of the end-to-end process that you're solving for or redesigning to be agentic."

Value is increasingly distributed across customer, operational, quality, productivity and risk outcomes. Not one response dominated when executives were asked to identify where agentic AI had delivered the most value. **Customer experience led by a small margin at 17%, followed by 12% in error reduction and quality improvement, and employee productivity, and compliance and risk management at 11%.** That mix suggests organizations are evaluating agentic AI against meaningful business outcomes as well as internal efficiency measures, rather than focusing solely on adoption or output.

Borrêga suggested that business value could be determined by connecting agentic work to a measurable business result and then optimizing the systems that support it. "We need to measure the outcome, which might be, for instance, building faster so we're able to put new products on the market faster, in which case we must be able to learn faster and react faster to competition. Then there's the question of how much risk we can control to make sure that the system operates as intended at scale," he said.

Conclusion:

From operational confidence to controlled capability

Executives reported that agentic AI had helped them most in responding to customer inquiries or service requests; resolving IT incidents and help desk support issues; processing finance, accounting or procurement tasks; and managing risk, compliance and security workflows. As agentic AI becomes embedded across more business processes, confidence will increasingly depend on what organizations can prove.

Governance policies remain mission critical. However, scaling agentic systems safely requires enterprises to translate those policies into operational controls that work consistently across agents, applications, data and models. Borrêga's closing advice was that organizations "shouldn't deploy AI without a clear business outcome. Instead, they should design the agentic process backward from that. That's where an end-to-end governance and control plane is critical."

That requires a four-part action framework:

- ✔ **Unify governance and orchestration across the ecosystem with a shared control plane.**
- ✔ **Contextualize and control by grounding your agentic workforce in a live operational context while enforcing security.**
- ✔ **Preserve full-stack choice across models, coding agents, deployment, hosting and architecture.**
- ✔ **Measure end-to-end outcomes, and design agentic processes backward from business value.**

In this case, "unified" means combining end-to-end control with outcome-based measurement. This is where leaders can create agentic systems more accountable and easier to adapt, optimize and scale as the business changes. Operational control is therefore not a brake on innovation but what lets enterprises change models, workflows and agents faster because they can do so with evidence and trust. As Harper said, "Clarity speeds organizations up. Ambiguity slows them down."

Methodology:

The findings of this research are based on an online survey conducted by CIO Dive from June 16 to June 29, 2026. A total of 858 executives from North America, Europe, Japan, India and Australia participated in the research. The sample was generated from CIO Dive databases.

The main roles participants played in the creation of custom software were defining application scope; evaluating development tools, vendors and/or partners; and defining business requirements.

About OutSystems

OutSystems is the Agentic Systems Platform built for the enterprise. Global organizations trust OutSystems to rapidly build mission-critical apps and agents, modernize legacy processes with agentic systems, and govern their entire AI portfolio across complex regulatory environments, all on a unified platform.

OutSystems is consistently recognized as a leader in enterprise software development and agentic enterprise orchestration by Gartner, IDC, and Forrester, and ranked #1 in Customer Satisfaction by users on G2. Business leaders, IT executives, and developers choose OutSystems to accelerate internal innovation without compromising reliability and security.

Founded in 2001, the OutSystems ecosystem includes more than 85 million end users, over 600 partners, and thousands of active customers in 75+ countries across 20+ industries. Learn more at www.outsystems.com.



Studio
by informa techtarget

Expert led.
Impact driven.

Studio is Informa TechTarget's global content studio offering brands an ROI rich tool kit: Deep industry expertise, first-party audience insights, an editorial approach to brand storytelling, and targeted distribution capabilities. Our trusted in-house content marketers help brands power insights-fueled content programs that nurture prospects and customers from discovery through to purchase, connecting brand to demand.

Learn more