UiPath 2026 AI and Agentic Automation Trends Report

Unlocking the Map



UiPath

© 2025 UiPath Inc. All rights reserved.
UiPath and the UiPath logo are trademarks or registered trademarks of UiPath Inc. in the United States and other countries.
All other trademarks mentioned are the property of their respective owners.

UiPath 2026 AI and Agentic Automation Trends Report

Unlocking the Map



Welcome to this year's **UiPath Trends Report**, our annual examination of the most powerful forces shaping the next wave of AI and agentic automation.

To formulate these trends, we have gathered perspectives from the UiPath automation ecosystem, which includes 10,000+ customers, 5,000+ partners, and 3 million UiPath Community members. We have also turned to our own experts: UiPath AI scientists, product and software teams, customer success teams, and sales and marketing experts. These diverse viewpoints, supplemented with a wide review of third-party research and analysis, have helped us to develop a practical and grounded view of where AI and agentic automation is headed, and what it will take to get there.

The theme of this year's trends is "Unlocking the Map." In gaming, this phrase describes how players gain access and visibility into hidden territories as they complete missions, discover new capabilities, and gain entry to previously inaccessible areas. What was once a shadowed outline becomes visible, navigable, and actionable.

And that's exactly what's happening in AI and agentic automation. Organizations now see the contours of what's possible and the possibilities contained within this new world. Boundaries and barriers—between data and action, between pilots and production, between initial investment and ultimate payoff—are falling. As more of the map unlocks, enterprises gain both vision and agency: the ability to navigate through a new agentic world with purpose, linking strategy, data, and technology into a cohesive system for growth.



This year, exploration gives way to execution as enterprises move from isolated trials toward connected, governed systems. Single-agent processes cede territory to multi-agent systems where agents, robots, and people work together across data, applications, and workflows. Organizations find a faster, smoother path to agentic ROI—leveraging the lessons learned in the last year, building out their infrastructure for scale, and speeding their way with prebuilt vertical solutions. They establish powerful, centralized orchestration and governance capabilities, and take steps to transform their data into competitive advantage.

2026 is the year when companies unlock the map that leads them into their agentic future. **Let's start the journey**.

Overview

UiPath 2026 AI and Agentic Automation Trends

1 **Necessity as the mother of reinvention**
AI's disruptions compel organizations to make bold changes in how they operate, compete, and allocate work.

2 **AI ROI at last**
Enterprises find their path from pilots to payout.

3 **Vertical ascent**
Focused agentic solutions take off and fly high.

4 **The power of the swarm**
Multi-agent systems land all over, transforming workflows for good...and for the better.

5 **Enter the command center**
Control moves to a new plane by centralizing orchestration, governance, and agent management.

6 **Gloves off, guardrails up**
Enterprises act decisively to ensure security, transparency, and control for every AI agent and agentic workflow.

7 **Data goes meta**
Companies double down on adding the context, structure, real-time access, and control that turn enterprise data into agent rocket fuel.



Trend #1

Necessity as the mother of reinvention

AI's disruptions compel organizations to make bold changes in how they operate, compete, and allocate work.

It's now clear that agent-centric operating models can dramatically outperform traditional ways of working—making it imperative for enterprises to reinvent themselves as agentic organizations. This requires adopting new operating systems built for orchestration, governance, and continuous optimization across an increasingly autonomous and interconnected digital enterprise.

Trend #1 Necessity as the mother of reinvention

Among executives, three-quarters predict agentic AI will reshape the workplace more profoundly than the Internet did, and 82% say it will transform their own industries within 18 months.[1] The vast majority have already brought some agents into their operations, and just about nine out of ten say they'll be increasing investment over 2026.[2] The momentum is inescapable. So is the need for profound reinvention.

While changes will ripple through leadership, talent, and strategy, the most striking area of transformation involves reconceiving the technology and operating systems that power the enterprise itself. It's become increasingly clear that operating models and technologies conceived for human workflows can't manage the scale, autonomy, and complexity of agentic environments.

Enterprises widely recognize this gap: in one study, 78% of C-suite executives agreed that gaining maximum benefit from agentic AI requires a new operating model built around agentic capabilities.[3]

This new agentic operating model requires new technologies that can orchestrate, govern, and continually optimize the work of agents, robots, and people across an increasingly complex digital ecosystem. In 2025, one-fifth of enterprises were already taking steps to rethink their operating models around an agentic core.[4] Throughout 2026, many more enterprises will take up the challenge of reinvention.

The new operating system will need to address several critical challenges, each with major implications for the technologies that support it.

1) Massive redistribution of work between humans and agents

The agentic era marks a radical redivision of labor between people and virtual workers. It's already happening: today, across roughly one-third of all occupations, AI systems perform more than a quarter of the tasks.[5] This disruption not only alters what people do and the skills they need, but also how work is visualized, assigned, and governed across the enterprise.

Technology implications: As people and virtual workers share more of the work, an orchestration layer that governs, monitors, and optimizes autonomous decisions in real time becomes paramount.

2) Agents moving into core, high-value workflows and processes

AI agents are extending automation into high-value, judgment-based processes such as decision making and risk management—creating new opportunities, but also new risks that call for strong oversight and control.[6]

Technology implications: Technology must support safe autonomy—automation that is auditable, explainable, and resilient, with human-in-the-loop capabilities. Enterprises will need embedded oversight capabilities

such as simulation, testing, and rollback, as well as continuous monitoring for compliance and accuracy. Reliability and assurance become as central to the operating system as speed and scale.

3) Continuous adaptation and real-time, flexible optimization

Because agents can learn, adapt, and self-optimize, change becomes continual as agentic systems fine-tune operations to maximize outcomes. This constant evolution challenges traditional ideas of stability and control.[7]

Technology implications: Operating systems will need to be modular, composable, and self-adaptive, integrating new tools, data sources, and agent capabilities without disruption. Enterprises will require simulation environments to test potential changes and real-time observability to track performance, safety, and alignment with business goals.

4) The blurring of boundaries

Value creation will increasingly happen across networks of interconnected agents spanning suppliers, partners, and customers. Think about procurement agents negotiating directly with vendor agents or consumer agents interacting with brand agents for personalized services. As these interactions multiply, the edges of the enterprise dissolve into a wider digital ecosystem.[8]

Technology implications: The operating system will need to be open, interoperable, and secure, enabling agents to operate across systems, vendors, and clouds. Shared protocols, APIs, and trust frameworks will be required to manage identity, permissions, and data exchanges.

The map ahead

In this new world, the OS becomes not just a control center but a bridge—connecting ecosystems while ensuring safety and transparency. As agentic capabilities scale, the question for enterprises will no longer be whether to adopt agents—but how to operate with them as core contributors.

Success will depend on implementing an operating system that integrates automation, governance, and orchestration into one adaptive capability. In 2026, necessity will indeed become the mother of reinvention—as enterprises reimagine not only their technology, but the very logic of how work gets done.

To-Do's for 2026

- **Frame your operating model**. Define the workflow between humans, robots, and agents—and required governance and orchestration.
- **Identify enabling capabilities**. Map the technologies and architecture your new operating model demands.
- **Set your build-out strategy**. Determine whether to build or buy, and identify external partners and platforms.
- **Move ahead to implement**. Act rapidly to capture agentic advantages.



78%

of C-suite executives agree that gaining maximum benefit from agentic AI requires a new operating model.

IBM Institute for Busines Value, *Atentic AI's Strategic Ascent*, 2025.

"[Winning companies] will abandon existing operating models for ones designed around autonomous decision making capabilities."

IBM Institute for Business Value, *Agentic AI's Strategic Ascent*, 2025.



Trend #2

AI ROI at last

Enterprises find their path from pilot to payout.

Organizations approach 2026 with momentum and a mandate: use the lessons of 2025 to make agentic programs deliver. Budgets are rising, confidence is high, and the experiences of 2025 have brought organizations far up the learning curve. As the focus this year shifts from experiment to execution, a new playbook for agentic payoff brings ROI within reach.

2025 laid the groundwork for the agentic era. Starting from near zero, enterprise experimentation with agents surged: by mid-year, roughly 65% of organizations were piloting or deploying agentic systems.[1] [2] Yet many have struggled to scale their AI efforts. According to an MIT study, only 5% of companies are realizing meaningful financial returns[3] and other studies report that 70–80% of agentic initiatives have failed to reach enterprise scale.[4] [5]

Even so, enterprises don't appear likely to curtail their efforts. Half of executives now rank agentic AI as their top AI investment priority for 2026.[6] IDC projects that agentic AI will represent 10–15% of IT spending in 2026 and grow at a 31.9% CAGR to 26% of budgets—about $1.3 trillion—by 2029.[7] In a world where agentic capability is the new foundation for competitiveness, it's clear that no one can afford to fall behind the curve.

2026: The year of "show me the money."

But while investment continues, there's an increasing focus on performance and impact. For instance, 80% of executives say their boards are asking for a defined strategy and demonstrable ROI for GenAI and agentic initiatives.[6] The mandate is clear: this year, agentic initiatives need to show proof of their scalability, efficiency, and business value.

The good news: many executives believe they can deliver. In a 3Q 2025 survey, 73% said their agentic initiatives would yield significant competitive advantage within a year, and 57% foresaw measurable ROI.[7]

The map ahead

As they move to capture more agentic value, executives will be following a new playbook for agentic payoff, which includes the following best practices:

- **Go where the pain pays.** Focus on high-pain, high-gain problems where agents can make a real difference—and where success builds momentum across the business.
- **Reinvent, don't retrofit.** Redesign processes from the ground up to take full advantage of what agents make possible.
- **Get the measure of your agents.** Old ROI metrics may not be the right ones. Focus on metrics that measure the real impact agents can deliver: business agility, better customer experience, and cost per transaction.
- **Find strength in swarms.** Multi-agent systems—fleets of specialized agents—are proving to be powerful in scaling performance.
- **Buy smart before you build**. Prebuilt, domain-specific agentic solutions and reusable workflows can deliver ROI faster and more predictably.

Trend #2 AI ROI at last

- **Don’t be late to orchestrate**. Orchestration synchronizes the handoffs among agents, robots, and humans across systems and departments.

- Keep control of the controls. Build in oversight, observability, and guardrails for every agent and workflow to sustain performance gains.

To-Do’s for 2026

- **Focus on high-ROI opportunities.** Learn how to identify them here: The Practitioner’s Guide to Agentic Automation.
- **Aim higher, win bigger.** Generate more ROI and impact by tackling a high-pain, high-gain process.
- **Set your agents up for success.** Learn how to maximize their impact with this report: Harness The Potential Of Agentic Automation At Your Organization.

73%

of executives predict their agentic projects will deliver value within 12 months.

Source: KPMG, *AI Quarterly Pulse Survey*, Q3 2025.

> **“Don’t think too small. If you focus on taking on the really tough problems to solve, the outcome can be much bigger.“**
>
> Daniel Dines, CEO, UiPath at FUSION, 2025



Trend #3

Vertical ascent

Focused agentic solutions take off and fly high.

This year, vertical agentic solutions continue their growth trajectory—covering more domains and enjoying broader adoption. Domain-tuned, tested, and integration-friendly, these configurations of agents, automation, models, and workflows offer faster deployment, measurable outcomes, and a reliable and effective pathway to scale.

Trend #3 Vertical ascent

In 2025, vertical agentic solutions—end-to-end, domain-specific offerings—began to gain real traction. In 2026, that momentum accelerates sharply, and by year's end, vertical agentic solutions will have firmly established their role in many enterprises' agentic portfolios.

Vertical solutions are poised for rapid growth because they meet enterprises' acute need to capture reliable and rapid ROI from their agentic initiatives. MIT research has found that externally sourced or partnership-based AI projects are twice as likely to achieve meaningful outcomes as internal builds.[1] By reducing cost-to-build, time to value, and performance risk, prebuilt solutions offer a powerful new option for organizations seeking to leverage AI and automation to achieve breakthrough productivity, resilience, and growth.

Solutions can be launched faster and deliver strong results because they include all the prebuilt elements and capabilities required for deployment, integration, and ongoing production:

- Specialized agents with predefined roles
- Optimized process design
- Workflow orchestration coordinating agents, people, applications, and systems
- Predefined data schema
- Fine-tuned, specialized models
- Industry-standard compliance frameworks including rules, escalation paths, and audit controls
- Built-in oversight and analytics for transparency, auditability, and performance tracking
- Integrations into enterprise ecosystems

The best vertical solutions are also highly configurable, for instance, making it possible for enterprises to use their own data and models. Workflows, decision frameworks, and guardrails can be adapted to enterprises' regulatory, security, and reporting standards, allowing companies to integrate solutions seamlessly into operations while maintaining compliance and control.

The map ahead

In 2026, vertical agentic solutions will become an essential part of many enterprises' AI strategies. Organizations will adopt a thoughtful mix of "make" and "buy," using vertical solutions to accelerate transformation in a select group of processes.

Attractive candidates for external solutions will be data-intensive, always-on processes where accuracy, decision speed, and throughput matter. They'll be processes where improvements can yield big business improvements, and where AI agents can deliver outsized impact.

Some examples of solutions already in market:

- Loan origination and compliance in financial services
- Claims and denials management in healthcare
- Supply chain management and optimization in manufacturing
- Inventory management, pricing, and merchandizing optimization in retail

Trend #3 Vertical ascent

Analysts characterize the solutions market as early but fast maturing,[2] as more providers and customers enter the space. In 2026, expect adoption to accelerate as enterprises experience the tangible benefits of vertical agentic solutions—faster deployment, lower integration risk, embedded compliance, and measurable ROI.

To-Do's for 2026

- **Set the portfolio mix.** Determine your "build or buy" strategy for key processes, identifying places where a vertical agentic solution can reliably deliver faster ROI than doing it internally.
- **Build an assessment framework**. As solutions offerings proliferate, make sure to establish clear criteria for evaluation.

2x

External solutions are twice as likely to deliver measurable results than internally built solutions

Source: MIT, *GenAI Divide: State of AI in Business 2025*, 2025.

"Industry-tailored solutions truly represent a fundamental shift in how organizations can leverage AI and automation to achieve breakthrough productivity and business outcomes."

Graham Sheldon, Chief Product Officer, UiPath



Trend #4

The power of the swarm

AI's multi-agent systems land all over, transforming workflows for good...and for the better.

This year, the enterprise focus changes from building single agents to launching multi-agent systems (MAS)—teams of agents working together. The shift is accelerating as organizations see how much better these systems can perform, and how much more they can do. Expect enterprises to shift their focus to MAS—and to invest in orchestration, governance, and skills to support these agent swarms.

Trend #4

The power of the swarm

A multi-agent system (MAS) is a network of autonomous agents that collaborate to achieve shared objectives. Each agent has a role and specific capabilities: one may plan, another retrieve data, another analyze, another act.

Leveraging parallel processing, redundant roles, and agent specialization—and coordinating activities through an orchestration layer that governs context, timing, and handoffs—MAS can effectively harness the power of agents to work together to execute end-to-end workflows. This approach can deliver substantial performance gains: up to 60% fewer errors, 40% faster execution, and 25% lower operating costs versus existing traditional processes.[1]

MAS also significantly extends the reach of agentic AI into processes where single agents have struggled, for example, workflows requiring diverse expertise, long horizons, and parallel steps. In one study, MAS-based approaches were able to complete complex processes 70% more often than single agents.[2]

The market is moving to embrace multi-agent systems. The number of enterprises moving beyond pilots is expected to double this year, and 75% of organizations plan to deploy multi-agent frameworks within the next 18 months.[3] Growth will continue: the global MAS market, which was $6.3 billion in 2025, is projected to grow at a 45.5% CAGR through 2034.[4]

Where the swarm is landing

Practitioners in many industries are successfully applying multi-agent systems to their complex, hard-to-automate processes.

Industry / Function	Processes	Key Drivers of Adoption
Banking & Insurance	KYC onboarding, fraud detection, claims adjudication, collections	Heavy compliance requirements, long multi-step workflows, need for speed + accuracy, regulatory pressure
Call Centers	Inquiry triage, resolution drafting, compliance/tone review	High cost of human labor, demand for faster response times, need to scale capacity without scaling headcount
IT & HR Service Operations	Case triage, approvals, knowledge retrieval, employee onboarding	High volume of repetitive service requests, pressure to reduce cost-to-serve, and improve employee experience
Professional Services	Contract review, audits, tax advisory, client onboarding	Knowledge-heavy tasks, need for explainability, client-facing risk if errors occur
Retail & Consumer	Dynamic pricing, demand forecasting, promotions, customer personalization	Competitive differentiation, margin pressure, demand for personalization at scale
Supply Chain & Logistics	Order-to-cash, network replanning, exception handling, demand/supply balancing	Urgent need for speed, cost efficiency, and resilience in volatile markets

The map ahead

As many first movers have discovered, implementing multi-agent systems isn't as simple as deciding to build their agentic automation and then launching it. You will likely need a significant amount of foundation-building.

As you plan and launch your MAS implementation in 2026, keep in mind these Requirements and Actions—best practices from UiPath agentic automation experts and UiPath customers.

Requirements	Actions
• Process intelligence: to help you understand current processes—and simulate and assess new approaches • Orchestration: to be able to assign tasks, track state, escalate exceptions, and manage agent teams. • Interoperability protocols: to enable agents from different vendors to collaborate. • Governance, safety frameworks, and capabilities: to reduce risks—policy-as-code, audit trails, explainability, and jailbreak protection • Security and compliance: frameworks and rules; visibility and monitoring capabilities • Job redesign and human-in-the-loop: to provide the structure, training, and tools to enable people to manage agents • KPIs and measurements: for people in new jobs, and for agents, too	• Ensure you have required capabilities and technology in process intelligence, orchestration, governance • Redesign design: rethink the design process, requirements, and goals • Set new KPIs: existing KPIs, like cost reduction and speed, may not be sufficient to assess how the agents themselves are truly performing; error detection, agent cooperation success, throughput, and resilience may be more helpful • Expand governance: add in new layer of runtime oversight and simulation testing before scale • Redefine jobs and build skills: shift roles toward monitoring, auditing, and optimizing multi-agent workflows

The power of the swarm

To-Do's for 2026

- **Engineer for orchestration.** Design an operating layer that manages agent coordination, permissions, and performance across workflows.
- **Build in governance.** Implement policy-as-code, audit trails, and observability frameworks to maintain trust and compliance.
- **Secure the swarm**. Integrate cybersecurity and access controls tailored to multi-agent architectures.
- **Build your skills in reskilling.** Give teams the new skills they need to manage hybrid human-agent ecosystems.

45%

of organizations scaling AI agents are piloting or scaling multi-agent systems

Source: Capgemini, *Harnessing the value of AI: Unlocking Scalable Advantage*, 2025.

"Risk increases sharply from single to multi-agent systems unless governance evolves in parallel."

Reid Blackman, CEO, Virtue Consultants, *Harvard Business Review*, June 2025.



Enter the command center

Organizations take control over agentic operations by centralizing orchestration, governance, and agent management.

As agentic automation extends across their core processes, enterprises move quickly to establish an operational infrastructure to unify oversight, compliance, control, and orchestration.

Trend #5 Enter the command center

It's become increasingly clear that organizations' adoption of AI in all its forms has advanced faster than their ability to govern, manage, and orchestrate it. Consider: in one study, nearly 75% of respondents had integrated AI into core operations—yet only one-third had mature governance controls in place.[1] Going into 2026, three-quarters of organizations were still 'in process' of developing their frameworks.[2]

For agentic AI specifically, the governance gap may be even wider: Everest Group estimates that only about 1% of firms have mature management infrastructures capable of orchestrating and governing agents effectively.[3]

This is a critical gap, and one that is getting wider as organizations deploy more agents and give them access to a wider range of processes. The adoption of multi-agent systems (MAS) adds yet more complexity, as these systems require sophisticated orchestration capabilities to direct, integrate, and monitor highly autonomous agents across systems, data, and workflows.

As organizations seek to close their governance gaps, it's become clear that traditional oversight models and methods—manual reviews, fragmented controls, and post-hoc audits—cannot provide the continuous, embedded visibility and real-time control that agentic environments demand. So many organizations are establishing a new operational layer—an agentic command center-—to centralize and integrate governance, control, and orchestration.

Meet the agentic command center

An enterprise control plane that embeds accountability and responsiveness into daily operations, the command center brings together the tools, data, and oversight needed to monitor, control, and coordinate agentic operations across the enterprise.

Core features include:

- **Centralized orchestration.** Agents, robots, and humans are coordinated across workflows and enterprise systems.
- **Governance baked in.** Policy-as-code, segregation of duties, and role-based access are all embedded at runtime.
- **Observability and resilience.** End-to-end tracing, cost accounting, failover mechanisms, and simulation sandboxes for testing agent plans are included.
- **Flexibility at scale.** Enables enterprises to bring in LLMs and regional models, third-party agents, and synthetic datasets while retaining oversight.

Enter the command center

What's in the center?

The agentic command center brings together an integrated set of capabilities to ensure full governance and orchestration for agentic operations.

Component	Purpose	Capabilities
Orchestration & workflow engine	To ensure coherent sequencing, routing, and coordination of agent tasks across domains	Task routing, dynamic chaining, retry/fallback logic, dependency resolution, scheduling
Governance & policy layer	To embed guardrails, compliance, and domain constraints	Policy-as-code enforcement, data access controls, role-based constraints, guardrails, rule versioning
Observability & audit/traceability	To provide transparency into decisions, behaviors, and performance	Logging, metrics, explainability, forensic traceability, anomaly detection
Lifecycle & change management	To control agent creation, deployment, updates, decommissioning	Versioning, testing/sandboxing, rollback, staging environments
Integration & extensibility fabric	To enable connections to external models, systems, and data	Adapter modules, plug-ins, model connectors, API gateways, data pipelines

The map ahead

In 2026, organizations will devote significant time, energy, and investment to ensuring their agentic operations have the governance frameworks and infrastructure needed to scale safely and effectively. Adoption of centralized approaches will continue to gain traction through 2026 and beyond. Analysts predict that by 2028, 70% of organizations deploying multi-agent and multi-LLM systems will use centralized orchestration platforms.[5]

To-Do's for 2026

- **Centralize control.** Move from using disparate, distributed tools toward establishing a centralized, integrated control layer that unifies oversight, governance, and orchestration across all agentic operations.

70%

of companies deploying multi-agent and multi-LLM systems will use centralized orchestration platforms by 2028

Source: Gartner, *Top Strategic Technology Trends for 2025: Agentic AI, 2024*.

“What distinguishes high-impact enterprise deployments of AI agents is the shift from many isolated pilots to industrial-scale orchestration and governance.”

McKinsey & Company, *Seizing the Agentic AI Advantage*, June 2025



Trend #6

Gloves off, guardrails up

Enterprise act decisively to ensure security, transparency, and control for every AI agent and agentic workflow.

As agents gain real autonomy—accessing data, making decisions, and executing actions—their security stakes skyrocket. In 2026, enterprises will turn from experimentation to enforcement, building 'trust by design' into every layer of the agentic stack. From governance-as-code and human-in-the-loop workflows to real-time observability, organizations are wiring in control systems to keep agents safe, compliant, and high-performing.

Trend #6 Gloves off, guardrails up

To harness AI agents at scale, enterprises must make them not only capable but trustable. That means building systems that ensure security, transparency, and control from the first line of code to the last workflow handoff.

Addressing this challenge is now a top enterprise priority. 96% of IT and security leaders view AI agents as a rising risk that must be addressed, and 92% agree that governing agents is essential. Yet fewer than half (44%) have formal policies in place.[1]

To ensure that agents 'do the right thing' every time, guidelines and guardrails should be directly built into them—codifying policies, permissions, and approval logic as part of their design and operation.[2]

This approach builds on long-standing software principles such as security by design, privacy by design, and DevSecOps, but reinterprets them for the agentic era.[3]

In 2026, embedding guardrails directly into agents will emerge as the leading approach to ensuring security and control. This approach will result in governance that extends across the entire agent lifecycle: control that is executed as code, validated by humans, and observable in real time. Organizations will also take advantage of the expanding range of technologies that make this possible, leveraging built-in policy engines, security frameworks, and orchestration features from leading platform providers to embed safeguards at the core of every agent.[4]

Technologies enabling the guardrails

The adoption of embedded governance across the agent lifecycle is being accelerated by a wave of technology investment and innovation.[5] Platform providers are building native capabilities that make it easier to hardwire safety, oversight, and control into agent design and operations.

These technologies allow enterprises to put lifecycle principles into practice—turning governance-as-code, human oversight, and observability into configurable system features rather than custom engineering challenges.

Technology companies are embedding governance and policy functions directly into their offerings—introducing built-in policy engines, approval frameworks, and audit controls that enable real-time supervision of agents in production. Providers of third-party or domain-specific agents are following suit, adding runtime governance features that ensure agents adhere to enterprise data policies, privacy requirements, and usage limits wherever they operate. Meanwhile, industry groups are developing interoperability protocols that enable enterprises to weave together offerings from multiple vendors seamlessly.[6]

Gloves off, guardrails up

Embedding governance across the agent lifecycle

DESIGN How coded		RUNTIME How deployed & constrained		ASSURANCE How observed & improved	
Governance-as-code	**Human-in-the-loop by design**	**Least-privilege access; isolation**	**Secure, data-native architecture**	**Continuous control & observation**	**Quality & adversarial testing**
Policies encoded directly into an agent's runtime	Human review woven into automated flows	Agents limited to strictly necessary tools & data	Agents run inside secure data planes	Action monitored in real time	Agents pressure-tested, constantly challenged
Executable rules define who and what an agent can access and when—ensuring approved actions, consistent behavior, and auditable decisions.	Agents propose, humans approve, and agents and robots execute steps—combining precision, accountability, oversight, and speed.	Segmented environments, temporary credentials, and scoped tool catalogs keep sensitive systems protected without slowing work.	Allows agents to automatically inherit masking, permissions, and lineage controls.	Instrumented logs, anomaly alerts, and operator kill switches provide transparency and responsiveness.	Red-team simulations and benchmarking against established industry standards reveal weaknesses early and harden agents against attack.
Enabling technologies					
Governance-aware agent builders Allow policies, approvals, and explainability to be designed directly into agents **Guardrails & orchestration frameworks** Set design-time and runtime configuration for safe agent behavior, topic control, and PII filtering		**Agent-native data platforms** Ensure secure data access and policy inheritance during agent operation **Agent observability suites** Deliver live monitoring, step-level tracing, and replayable logs to ensure compliance and responsiveness **Process orchestration systems** Coordinate the execution of agents, people, and robots across enterprise systems and workflows		**Interoperability protocols** Enable secure, policy-compliant collaboration among agents operating across different platforms or vendors **Quality and adversarial testing** Help identify vulnerabilities and strengthen agent performance over time	

The map ahead

In 2026, enterprises will focus on formalizing policies for agent lifecycle management—covering approvals, data permissions, and accountability—and make governance-as-code a top initiative. Automation, security, and data teams will align around shared frameworks that keep agents compliant, explainable, and secure as they scale.

Ensuring agent performance will also become a priority, with continuous observability, testing, and model evaluation baked into the agent management framework from inception.

To-Do's for 2026

- **Set agent rules and permissions.** Codify the guides agent behavior and automate approvals.
- **Build in improvement.** Set up continuous agent monitoring and testing.
- **Align people and platforms.** Make sure governance, data, and automation efforts advance together.

96%

of IT and security leaders view AI agents as a rising risk that must be addressed

Source: SailPoint, *The Rising Risk of AI Agents: Expanding the Attack Surface*, 2025.

> **“How do you make sure someone isn’t attacking the agents or extracting information they shouldn’t? Security has to be baked in from the start.”**
>
> Maxim Ioffe, Global Intelligent Automation Leader, Wesco at UiPath FUSION 2025



Trend #7

Data goes meta

Companies double down on adding the context, structure, real-time access, and control that turn enterprise data into agent rocket fuel.

As enterprises expand their use of GenAI and agentic AI, data quality, structure, and context are becoming decisive advantages. In 2026, organizations will focus on enriching and governing their data—building real-time, trusted, and semantically rich systems that give AI agents the understanding they need to perform with accuracy, confidence, and control.

Trend #7 Data goes meta

Managing enterprise data has long been a defining challenge—and a source of opportunity—for organizations. As they expand their use of GenAI and agentic AI, that opportunity becomes even more strategic. AI agents and GenAI models depend on a reliable flow of high-quality, trustable data (often in real time) to perform at their best.[1]

Addressing data quality and usability is a top concern for leaders as they seek to scale their agentic initiatives. In a recent survey, 82% of executives named 'organizational data quality' as the greatest barrier to achieving their GenAI goals—up from 56% just six months earlier.[2] A strong data foundation can spell the difference between AI initiatives that contribute significant value and those that deliver only incremental gains. McKinsey analysis shows that companies with mature data capabilities are three times more likely than peers to capture at least 20% of their EBIT from data and analytics initiatives.[3]

In 2026, organizations will direct new energy and investment toward strengthening this foundation. Five areas they'll be focusing on:

1) Giving data meaning: metadata and ontologies

Access alone isn't enough—agents also need to understand what the data means. That understanding comes from metadata and ontologies. Metadata describes datasets: what they contain, who owns them, and how reliable they are. Ontologies go further, defining how things connect—a policy belongs to a customer; a claim settles a policy; a payment closes a claim. Together, they give AI systems a structured view of the enterprise world.

When data is enriched with these semantic layers, AI models perform dramatically better. For example, in one study, ontology-enriched data lifted a large language model's accuracy from 16% to 54%.[4]

2) Providing the right data to agents

AI agents are only as effective as the data they can reach. To reason and act reliably, they need access to real-time, trustworthy, and governed information. Enterprises are investing in data fabrics and modernized architectures that connect sources virtually—CRM, ERP, warehouses, and unstructured content—through shared governance and policy enforcement. These systems allow agents to query and act on live data in place while automatically applying lineage, permissions, and security controls. In 2026, more organizations will adopt zero-copy, governed architectures, ensuring agents always operate on current, compliant data.

3) Coding in governance

As AI agents act with greater autonomy, governance is moving from policy to practice—embedded directly into data systems. Policy-as-code allows business rules, access controls, and compliance requirements to travel with the data itself. Every query or action an agent takes is checked in real time against these encoded rules, ensuring safety, transparency, and auditability. Gartner projects that by 2028, 90% of enterprise AI systems will include real-time policy enforcement and observability frameworks to ensure trustworthy agent behavior.[5] This evolution transforms governance into an enabler of safe autonomy.

4) Creating an edge with real-time context

AI's value increasingly depends on how quickly it can sense and respond. In 2026, leading enterprises will expand event-driven architectures that feed agents continuous, real-time context—transactions, sensor data, customer actions. These systems give agents the situational awareness to make immediate, accurate decisions. Researchers report that organizations combining real-time data with intelligent automation achieve 25% faster decision cycles and 40% lower error rates than those relying on static datasets.[6, 7] For AI-driven operations, live context is becoming a critical differentiator—enabling in-the-moment precision.

5) Creating a moat with proprietary data

Public data trains models to understand the world; proprietary data trains them to understand the enterprise. Customer histories, operational telemetry, service logs, and other unique data reflect each organization's distinct workflows, logic, and customer interactions. This data enables agents to act contextually, make consistent and compliant decisions, anticipate needs, and personalize at scale.

At the individual model level, introducing proprietary data has been shown to reduce error rates by up to 40%.[8] At the enterprise level, organizations that integrate proprietary data into their AI systems consistently outperform peers—one study found their EBITDA to be 25% higher.[9]

In 2026, that flywheel will become the true moat, powered by enterprises that turn governed, contextual data into a lasting engine of intelligence and advantage.

The map ahead

2026 will be the meta year—when enterprises create a living framework of meaning and control for their data. By enriching their data with structure, semantics, and governance, organizations will give AI systems what they need most: understanding. That understanding is the bridge to the next phase—data-native agents—AI that lives inside data, reasoning and acting with the same confidence and context as the humans it assists.

To-Do's for 2026

- **Make data agent-ready.** Connect data and ensure governance, quality, and access.
- **Set the context.** Enrich data with metadata, ontologi.es, and policy-as-code rules.
- **Be proprietary.** Identify high-value enterprise data and make sure it's structured, governed, and ready for real-time use.

393.9 ZB

the amount of data that will be created, captured, replicated, and consumed worldwide by 2028

Source: IDC Global DataSphere

“Agents can only make the right decisions when they’re grounded in the right context. Enterprise data today is messy and fragmented. When you organize that data, agents can reason over it accurately and deliver reliable results at scale.”

Jerry Liu, Founder & CEO, LlamaIndex at UiPath FUSION, 2025

Sources

Trend #1

Necessity as the mother of reinvention.

1. PwC, *AI Agent Survey*, 2025.
2. Ibid.
3. Gartner*, *Top Strategic Technology Trends*, 2025.
4. IBM Institute for Business Value, *Agentic AI's Strategic Ascent*, 2025.
5. Stanford Digital Economy Lab, *Future of Work with AI Agents*, 2025.
6. IDC, *Future of Digital Business*, 2025.
7. Everest Group, *A Practitioner's Guide to Agentic Automation*, 2025.
8. McKinsey, *Seizing the Agentic AI Advantage*, 2025.

Trend #2

AI ROI at last.

1. PwC, *AI Agent Survey,* May 2025.
2. McKinsey & Company, *Seizing the Agentic AI Advantage: A CEO Playbook*, June 2025.
3. MIT Sloan Management Review, *The State of AI 2025*, October 2025.
4. Accenture, *Front-Runners' Guide to Scaling AI*, 2025.
5. Wipro, *State of Data4AI Report 2025*, 2025.
6. KPMG, *Executive Pulse Q3 2025: GenAI Strategy and Investment Trends,* 2025.
7. IDC news release, *"Agentic AI to Dominate IT Budget Expansion Over Next Five Years,"* Aug.26, 2025.

Trend #3

Vertical ascent.

1. Everest Group, *Innovation Watch: Agentic AI Products*, 2025
2. Massachusetts Institute of Technology *GenAI Divide: State of AI in Business*, 2025

Trend #4

The power of the swarm.

1. Everest Group, *A Practitioner's Guide to Agentic Automation*, August 2025.
2. Ibid.
3. PwC. *AI Agent Survey*. May 2025.
4. Gartner*, *Emerging Tech Disruptors*, August 2025.

**GARTNER is a registered trademark and service mark of Gartner, Inc. and/or its affiliates in the U.S. and internationally. All rights reserved.*

Trend #5

Enter the command center.

1. McKinsey & Company,*The State of AI: How Organizations Are Rewiring to Capture Value*, March 2025.
2. PwC, *AI Agent Survey*, May 2025.
3. Everest Group, *A Practitioner's Guide to Agentic Automation*, August 2025.
4. Gartner*, *Predicts 2025: AI Risks and Security Implications of Agent Sprawl*, October 2024.
5. Gartner*, *Top Strategic Technology Trends for 2025: Agentic AI,* October 2024.

Trend #6

Gloves off, guardrails up.

1. PwC, *AI Agent Survey*. May 2025.
2. Everest Group, *A Practitioner's Guide to Agentic Automation*, August 2025.
3. Gartner*, *Predicts 2025: AI Risks and Security Implications of Agent Sprawl*, October 2024.
4. Gartner*, *Top Strategic Technology Trends for 2025: Agentic AI*, October 2024.
5. IDC, *Worldwide AI and Automation Infrastructure Forecast, 2025–2029*, January 2025.
6. McKinsey & Company, *Seizing the Agentic AI Advantage*, June 2025.

Trend #7

Data goes meta.

1. McKinsey & Company, *Seizing the Agentic AI Advantage*, 2025.
2. KPMG, *Executive Pulse Q3 2025: GenAI Strategy and Investment Trends*, 2025.
3. McKinsey & Company, *The Future of AI in the Insurance Industry*, 2025.
4. Kayali, Moe, et al., *"Mind the Data Gap: Bridging LLMs to Enterprise Data Integration,"* arXiv preprint, 2025.
5. Gartner*, *Top Strategic Technology Trends for 2025: Agentic AI*, 2025.
6. McKinsey & Company, Op. cit, #1
7. McKinsey & Company, Op. cit. #3
8. McKinsey & Company, Op. cit. #1
9. McKinsey & Company, Op. cit. #3

Fair Harbor

Fair Harbor notification

Statements we make in this presentation may include statements which are not historical facts and are considered forward-looking within the meaning of the Private Securities Litigation Reform Act of 1995, which are usually identified by the use of words such as “anticipates,” “believes,” “estimates,” “expects,” “intends,” “may,” “plans,” “possible,” “projects,” “outlook,” “seeks,” “should,” “will,” and variations of such words or similar expressions, including the negatives of these words or similar expressions.

We intend these forward-looking statements to be covered by the safe harbor provisions for forward-looking statements contained in Section 27A of the Securities Act of 1933, as amended, and Section 21E of the Securities Exchange Act of 1934, as amended, and are making this statement for purposes of complying with those safe harbor provisions.

These forward-looking statements include, but are not limited to, statements regarding: our ability to drive and accelerate future growth and operational efficiency and grow our platform, product offerings, and market opportunity; our business strategy; plans and objectives of management for future operations; the estimated addressable market opportunity for our platform and the growth of the enterprise automation market; the success of our platform and new releases including the incorporation of AI; the success of our collaborations with third parties; our customers’ behaviors and potential automation spend; and details of UiPath’s stock repurchase program. Forward-looking statements involve known and unknown risks, uncertainties, and other factors that may cause our actual results, performance, or achievements to be materially different from any future results, performance, or achievements expressed or implied by the forward-looking statements. These risks include, but are not limited to, risks and uncertainties related to: our expectations regarding our revenue, annualized renewal run-rate (ARR), expenses, and other operating results; our ability to effectively manage our growth and achieve or sustain profitability; our ability to acquire new customers and successfully retain existing customers; the ability of the UiPath Platform™ to satisfy and adapt to customer demands and our ability to increase its adoption; our ability to grow our platform and release new functionality in a timely manner; future investments in our business, our anticipated capital expenditures, and our estimates regarding our capital requirements; the costs and success of our marketing efforts and our ability to evolve and enhance our brand; our growth strategies; the estimated addressable market opportunity for our platform and for automation in general; our reliance on key personnel and our ability to attract, integrate, and retain highly-qualified personnel and execute management transitions; our ability to obtain, maintain, and enforce our intellectual property rights and any costs associated therewith; the effect of significant events with macroeconomic impacts, including but not limited to military conflicts and other changes in geopolitical relationships and inflationary cost trends, on our business, industry, and the global economy; our reliance on third-party providers of cloud-based infrastructure; our ability to compete effectively with existing competitors and new market entrants, including new, potentially disruptive technologies; the size and growth rates of the markets in which we compete; and the price volatility of our Class A common stock.
Further information on risks that could cause actual results to differ materially from our guidance and other forward-looking statements can be found in our Annual Report on Form 10-K for the fiscal year ended January 31, 2025 filed with the United States Securities and Exchange Commission (SEC), in our Quarterly Reports on Form 10-Q filed with the SEC, and in other filings and reports that we may file from time to time with the SEC. Any forward-looking statements contained in this presentation are based on assumptions that we believe to be reasonable as of this date. Except as required by law, we assume no obligation to update these forward-looking statements.
Our fiscal year end is January 31, and our fiscal quarters end on April 30, July 31, and October 31. All third-party trademarks, including names, logos and brands, referenced by us in this presentation are property of their respective owners. All references to third-party trademarks are for identification purposes only. Such use should not be construed as an endorsement of the products or services of us.