

AI Agents Step Up: From Chat to Action

Executive Summary

AI agents have evolved from being mere chatbots into autonomous team members tackling real work. In the last two days, new releases and enterprise moves underscore the pace of change — and highlight new demands for governance and trust.

AI Agents Become Autonomous Team Players

In one dramatic example, AI lab Anthropic this week launched Claude Tag, a persistent AI teammate integrated into Slack for its enterprise customers (venturebeat.com [1]) (venturebeat.com [2]). Unlike a traditional chatbot that only responds to one user at a time, Claude Tag can be invoked by anyone in a channel and maintain shared context as it interacts with multiple team members. Colleagues can delegate coding, research, or data analysis tasks to the AI by tagging @Claude in chat, and it will break the request into stages, execute each step with the tools it has access to, and return results along with an audit trail of its actions (venturebeat.com [3]) (venturebeat.com [4]).

This new AI agent doesn't just wait for instructions – it can also proactively assist the team. With an ambient mode enabled, Claude Tag monitors project discussions to surface relevant information from across the organization and even follows up on tasks that have stalled, nudging team members or autonomously continuing the work until it's completed (venturebeat.com [5]). It operates asynchronously, meaning it can pursue long-running projects over hours or days without constant human intervention. In effect, Anthropic has created a digital colleague that learns the business context, takes initiative, and collaborates with humans in real time.

The step-change in capability is striking. Anthropic reports that 65% of its own product team's code is now generated by its internal version of Claude Tag (venturebeat.com [6])— a dramatic illustration of how quickly an AI agent can ramp up productivity once it masters a team's workflows and knowledge base. And it's not just Anthropic pushing in this direction. Google has been upgrading its NotebookLM system to autonomously gather sources, create diagrams, and perform multi-step research tasks, moving it closer to a true AI research assistant (imfounder.com [7]). Microsoft, meanwhile, is enhancing its Copilot ecosystem to automate more complex, multi-application workflows. The clear trend is that AI is evolving from merely answering questions to directly carrying out work. This opens new opportunities to streamline operations, but it also raises the stakes for managing how these digital "teammates" behave and are supervised.

References:

- [1] venturebeat.com — <https://venturebeat.com/technology/anthropic-launches-claude-tag-replacing-its-slack-app-with-a-persistent-ai-teammate-that-learns-monitors-and-works-autonomously#:~:text=teammate%20that%20learns%2C%20monitors%20and,existing%20Claude%20in%20Slack%20app>
- [2] venturebeat.com — <https://venturebeat.com/technology/anthropic-launches-claude-tag-replacing-its-slack-app-with-a-persistent-ai-teammate-that-learns-monitors-and-works-autonomously#:~:text=technology%20leaders%20who%20have%20spent,code%20is%20now%20created%20by>
- [3] venturebeat.com — <https://venturebeat.com/technology/anthropic-launches-claude-tag-replacing-its-slack-app-with-a-persistent-ai-teammate-that-learns-monitors-and-works-autonomously#:~:text=flows%20through%20the%20tool%20it,The>

- [4] [venturebeat.com — https://venturebeat.com/technology/anthropic-launches-claude-tag-replacing-its-slack-app-with-a-persistent-ai-teammate-that-learns-monitors-and-works-autonomously#:~:text=context%20from%20other%20Slack%20channels,spend%20much%20more%20of%20our](https://venturebeat.com/technology/anthropic-launches-claude-tag-replacing-its-slack-app-with-a-persistent-ai-teammate-that-learns-monitors-and-works-autonomously#:~:text=context%20from%20other%20Slack%20channels,spend%20much%20more%20of%20our)
- [5] [venturebeat.com — https://venturebeat.com/technology/anthropic-launches-claude-tag-replacing-its-slack-app-with-a-persistent-ai-teammate-that-learns-monitors-and-works-autonomously#:~:text=context%20from%20other%20Slack%20channels,spend%20much%20more%20of%20our](https://venturebeat.com/technology/anthropic-launches-claude-tag-replacing-its-slack-app-with-a-persistent-ai-teammate-that-learns-monitors-and-works-autonomously#:~:text=context%20from%20other%20Slack%20channels,spend%20much%20more%20of%20our)
- [6] [venturebeat.com — https://venturebeat.com/technology/anthropic-launches-claude-tag-replacing-its-slack-app-with-a-persistent-ai-teammate-that-learns-monitors-and-works-autonomously#:~:text=implications%20for%20enterprise%20workflow%2C%20governance%2C,How%20Claude%20Tag%20works](https://venturebeat.com/technology/anthropic-launches-claude-tag-replacing-its-slack-app-with-a-persistent-ai-teammate-that-learns-monitors-and-works-autonomously#:~:text=implications%20for%20enterprise%20workflow%2C%20governance%2C,How%20Claude%20Tag%20works)
- [7] [imfounder.com — https://imfounder.com/science-tech/ai/ai-updates-june-2026-siri-gemini-claude/#:~:text=creation%20Product%20redesigns%2013,14](https://imfounder.com/science-tech/ai/ai-updates-june-2026-siri-gemini-claude/#:~:text=creation%20Product%20redesigns%2013,14)

From Chat to Full Service in Financial Firms

The financial services sector is likewise pushing past the era of basic Q&A bots toward more capable AI agents. Digital banking platform Backbase, for example, announced it will acquire Kasisto, a company known for powering conversational banking apps (pulse2.com [1]). The goal is to move beyond today’s isolated chatbots to what Backbase calls “governed AI agents” that can actually resolve customers’ requests across various channels and back-end systems (pulse2.com [2]). Many banks’ current AI assistants only answer simple questions — for anything complex, the interaction still ends with a human agent finishing the job (pulse2.com [3]). By embedding Kasisto’s technology, Backbase plans to let banks deploy virtual agents that manage the entire “intent-to-resolution” process: not just telling a customer their balance or locking a lost card, but executing the transfer, blocking the card, initiating follow-ups, and updating all internal records along the way.

Kasisto’s platform brings deep banking domain knowledge and pre-built regulatory guardrails to support this vision (pulse2.com [4]). Backbase’s CEO described the combined solution as a “unified frontline,” where customers, employees, and AI agents share one integrated platform and a single source of truth about each client interaction (pulse2.com [5]). In practice, that means an AI agent can tap into core banking systems and workflows, so a request made over a chat channel can trigger actions in contact center tools, databases, or even third-party services — all orchestrated seamlessly by the AI (pulse2.com [6]). Importantly, these autonomous actions are constrained by business rules and oversight: the “governed” in governed AI agent. Every automated step, from verifying a customer’s identity to executing a transaction, can be monitored and audited.

Pacesetters like Backbase see such agentic automation as a competitive differentiator. As CEO Jouk Pleiter put it, the acquisition moves the firm decisively into the era of customers expressing intent naturally and the bank fulfilling it through intelligent, automated execution (pulse2.com [7]). For bank executives, this development demonstrates that AI isn’t just a flashy add-on for customer service — it’s becoming a core part of the operating fabric. The challenge now will be ensuring these digital agents truly deliver on their promise: simplifying customer journeys without introducing new risks. That will require tight integration with legacy systems, extensive testing, and robust oversight to maintain trust with customers and regulators.

References:

- [1] [pulse2.com — https://pulse2.com/backbase-acquires-kasisto-to-advance-agentic-banking/#:~:text=To%20Advance%20Agentic%20Banking%20By,use%20AI%20agents%20that%20answer](https://pulse2.com/backbase-acquires-kasisto-to-advance-agentic-banking/#:~:text=To%20Advance%20Agentic%20Banking%20By,use%20AI%20agents%20that%20answer)
- [2] [pulse2.com — https://pulse2.com/backbase-acquires-kasisto-to-advance-agentic-banking/#:~:text=acquisition%20is%20designed%20to%20help,strategy%2C%20which%20is%20focused%20on](https://pulse2.com/backbase-acquires-kasisto-to-advance-agentic-banking/#:~:text=acquisition%20is%20designed%20to%20help,strategy%2C%20which%20is%20focused%20on)
- [3] [pulse2.com — https://pulse2.com/backbase-acquires-kasisto-to-advance-agentic-banking/#:~:text=acquisition%20is%20designed%20to%20help,and%20AI%20agents%20work%20with](https://pulse2.com/backbase-acquires-kasisto-to-advance-agentic-banking/#:~:text=acquisition%20is%20designed%20to%20help,and%20AI%20agents%20work%20with)
- [4] [pulse2.com — https://pulse2.com/backbase-acquires-kasisto-to-advance-agentic-banking/#:~:text=Kasisto%20brings%20banking,back%20resolution%20for%20customer](https://pulse2.com/backbase-acquires-kasisto-to-advance-agentic-banking/#:~:text=Kasisto%20brings%20banking,back%20resolution%20for%20customer)
- [5] [pulse2.com — https://pulse2.com/backbase-acquires-kasisto-to-advance-agentic-banking/#:~:text=acquisition%20supports%20Backbase%E2%80%99s%20Unified%20Frontline,service%2C%20contact%20centers](https://pulse2.com/backbase-acquires-kasisto-to-advance-agentic-banking/#:~:text=acquisition%20supports%20Backbase%E2%80%99s%20Unified%20Frontline,service%2C%20contact%20centers)
- [6] [pulse2.com — https://pulse2.com/backbase-acquires-kasisto-to-advance-agentic-banking/#:~:text=shared%20context%2C%20governed%20authority%2C%20and,The%20acquisition%20also](https://pulse2.com/backbase-acquires-kasisto-to-advance-agentic-banking/#:~:text=shared%20context%2C%20governed%20authority%2C%20and,The%20acquisition%20also)
- [7] [pulse2.com — https://pulse2.com/backbase-acquires-kasisto-to-advance-agentic-banking/#:~:text=future%20customers,built%20agentic%20AI%20and](https://pulse2.com/backbase-acquires-kasisto-to-advance-agentic-banking/#:~:text=future%20customers,built%20agentic%20AI%20and)

Built for the Enterprise: Governance and Context

In other industries, leaders are similarly pairing AI ambitions with domain expertise and control. This week, Seattle-based HR tech firm Syndio announced its first-ever acquisition: Embrace.ai, an “agentic AI” startup built to deploy automation across business workflows with a focus on governance and explainability ([www.geekwire.com \[1\]](https://www.geekwire.com/2026/syndio-bets-on-agentic-ai-with-first-acquisition-in-seattle-pay-equity-startups-history/#:~:text=history%20%E2%80%93%20GeekWire%20Syndio%20bets,to)) ([www.geekwire.com \[2\]](https://www.geekwire.com/2026/syndio-bets-on-agentic-ai-with-first-acquisition-in-seattle-pay-equity-startups-history/#:~:text=enterprise%20human%20capital%20management%20giant,powered%20compensation%20platform)). By bringing aboard an entire team of AI specialists, Syndio aims to accelerate its “Decision Intelligence for Pay” initiative — tools that help large employers analyze, govern, and improve pay and hiring decisions at scale ([www.geekwire.com \[3\]](https://www.geekwire.com/2026/syndio-bets-on-agentic-ai-with-first-acquisition-in-seattle-pay-equity-startups-history/#:~:text=revealed,the%20next%20generation%20of%20pay)). Embrace.ai’s co-founders (both Workday alumni) spent years implementing AI-driven processes inside enterprises, and their experience will help Syndio infuse its platform with safe, context-aware automation for complex HR decisions.

Syndio’s CEO Maria Colacurcio has emphasized that high-stakes processes like compensation require AI that deeply understands company-specific data, policies, and constraints ([www.geekwire.com \[4\]](https://www.geekwire.com/2026/syndio-bets-on-agentic-ai-with-first-acquisition-in-seattle-pay-equity-startups-history/#:~:text=history%20%E2%80%93%20GeekWire%20Syndio%20bets,to)). That’s why a governance-first, domain-trained approach is crucial. The Embrace.ai team was explicitly built around the idea that enterprise AI must be deployed with strong guardrails and clear explanations for its actions ([www.geekwire.com \[5\]](https://www.geekwire.com/2026/syndio-bets-on-agentic-ai-with-first-acquisition-in-seattle-pay-equity-startups-history/#:~:text=enterprise%20human%20capital%20management%20giant,powered%20compensation%20platform)). As one Embrace founder put it, an AI decision-maker needs to “support, not replace, human judgment” if it is to be trusted in sensitive matters of pay and equity ([www.geekwire.com \[6\]](https://www.geekwire.com/2026/syndio-bets-on-agentic-ai-with-first-acquisition-in-seattle-pay-equity-startups-history/#:~:text=revealed,the%20next%20generation%20of%20pay)). In practice, this means any recommendation an AI agent provides to a manager about salaries or promotions should come with transparent reasoning and the ability to defer to a human when uncertainty arises.

Beyond acquisition, investors are also funding the infrastructure needed for agentic AI. Consider the challenge of information retrieval: standard search engines return a list of links for humans to click, but AI agents need to ingest relevant data directly. Recognizing this gap, a startup called Seltz just raised \$12.5 million to build a specialized web search platform for AI agents ([siliconangle.com \[7\]](https://siliconangle.com/2026/seltz-raises-12-5-million-to-build-ai-search-platform)). Traditional search engines weren’t designed for the way AI systems query and read the web, which can involve dozens of long, detailed queries issued in parallel and looking for specific snippets in text, tables, or images that might never surface to a human user ([finance.yahoo.com \[8\]](https://finance.yahoo.com/2026/seltz-raises-12-5-million-to-build-ai-search-platform)). Seltz is one of several players in this space — others have already raised orders of magnitude more funding (one rival secured \$100 million at a \$2 billion valuation, and another raised \$250 million last month) to build “agentic search” solutions ([siliconangle.com \[9\]](https://siliconangle.com/2026/seltz-raises-12-5-million-to-build-ai-search-platform)). This surge of investment signals just how key supporting technologies like data pipelines, orchestration frameworks, and secure integration layers are becoming in the autonomous workflow era.

The thread running through all these developments is that generic AI tools are not enough. Whether in finance, HR, customer service, or other domains, effective AI agents need rich contextual understanding and strong oversight baked in from the start. Enterprises are learning that successful, scalable automation depends on sound architecture — unified data, clear ontologies, and continuous monitoring — to prevent the very issues that can derail AI initiatives. Without those foundations, as one analysis put it, organizations often “rush toward agentic AI without addressing the structural requirements,” and promising pilots collapse when exposed to the messy complexity of real business operations ([undercodenews.com \[10\]](https://undercodenews.com/2026/syndio-bets-on-agentic-ai-with-first-acquisition-in-seattle-pay-equity-startups-history/#:~:text=history%20%E2%80%93%20GeekWire%20Syndio%20bets,to)) ([undercodenews.com \[11\]](https://undercodenews.com/2026/syndio-bets-on-agentic-ai-with-first-acquisition-in-seattle-pay-equity-startups-history/#:~:text=enterprise%20human%20capital%20management%20giant,powered%20compensation%20platform)).

References:

- [1] [www.geekwire.com — https://www.geekwire.com/2026/syndio-bets-on-agentic-ai-with-first-acquisition-in-seattle-pay-equity-startups-history/#:~:text=history%20%E2%80%93%20GeekWire%20Syndio%20bets,to](https://www.geekwire.com/2026/syndio-bets-on-agentic-ai-with-first-acquisition-in-seattle-pay-equity-startups-history/#:~:text=history%20%E2%80%93%20GeekWire%20Syndio%20bets,to)
- [2] [www.geekwire.com — https://www.geekwire.com/2026/syndio-bets-on-agentic-ai-with-first-acquisition-in-seattle-pay-equity-startups-history/#:~:text=enterprise%20human%20capital%20management%20giant,powered%20compensation%20platform](https://www.geekwire.com/2026/syndio-bets-on-agentic-ai-with-first-acquisition-in-seattle-pay-equity-startups-history/#:~:text=enterprise%20human%20capital%20management%20giant,powered%20compensation%20platform)
- [3] [www.geekwire.com — https://www.geekwire.com/2026/syndio-bets-on-agentic-ai-with-first-acquisition-in-seattle-pay-equity-startups-history/#:~:text=revealed,the%20next%20generation%20of%20pay](https://www.geekwire.com/2026/syndio-bets-on-agentic-ai-with-first-acquisition-in-seattle-pay-equity-startups-history/#:~:text=revealed,the%20next%20generation%20of%20pay)
- [4] [www.geekwire.com — https://www.geekwire.com/2026/syndio-bets-on-agentic-ai-with-first-acquisition-in-seattle-pay-equity-startups-history/#:~:text=history%20%E2%80%93%20GeekWire%20Syndio%20bets,to](https://www.geekwire.com/2026/syndio-bets-on-agentic-ai-with-first-acquisition-in-seattle-pay-equity-startups-history/#:~:text=history%20%E2%80%93%20GeekWire%20Syndio%20bets,to)

startups-history/#:~:text=reporting%20into%20what%20it%20calls,the%20next%20generation%20of%20pay

[5] [www.geekwire.com](https://www.geekwire.com/2026/syndio-bets-on-agentic-ai-with-first-acquisition-in-seattle-pay-equity-startups-history/) — <https://www.geekwire.com/2026/syndio-bets-on-agentic-ai-with-first-acquisition-in-seattle-pay-equity-startups-history/#:~:text=an%20agent%20AI%20startup%20whose,for%20Pay%E2%80%9D%20%E2%80%94%20helping%20companies>

[6] [www.geekwire.com](https://www.geekwire.com/2026/syndio-bets-on-agentic-ai-with-first-acquisition-in-seattle-pay-equity-startups-history/#:~:text=strategy,powered%20compensation%20platform) — <https://www.geekwire.com/2026/syndio-bets-on-agentic-ai-with-first-acquisition-in-seattle-pay-equity-startups-history/#:~:text=strategy,powered%20compensation%20platform>

[7] [siliconangle.com](https://siliconangle.com/2026/06/24/agent-ai-infrastructure-startup-seltz-raises-12-5m-help-ai-agents-search-web-answers/#:~:text=web%20for%20answers%20,startup%20was%20founded%20by%20its) — <https://siliconangle.com/2026/06/24/agent-ai-infrastructure-startup-seltz-raises-12-5m-help-ai-agents-search-web-answers/#:~:text=web%20for%20answers%20,startup%20was%20founded%20by%20its>

[8] [finance.yahoo.com](https://finance.yahoo.com/technology/ai/articles/exclusive-seltz-startup-rebuilding-search-110000196.html#:~:text=founder%20and%20CEO%20Antonio%20Mallia,of%20representation%20that%20can%20be) — <https://finance.yahoo.com/technology/ai/articles/exclusive-seltz-startup-rebuilding-search-110000196.html#:~:text=founder%20and%20CEO%20Antonio%20Mallia,of%20representation%20that%20can%20be>

[9] [siliconangle.com](https://siliconangle.com/2026/06/24/agent-ai-infrastructure-startup-seltz-raises-12-5m-help-ai-agents-search-web-answers/#:~:text=extracting%20the%20exact%20text%2C%20table,Mallia%E2%80%99s%20company%20is%20also) — <https://siliconangle.com/2026/06/24/agent-ai-infrastructure-startup-seltz-raises-12-5m-help-ai-agents-search-web-answers/#:~:text=extracting%20the%20exact%20text%2C%20table,Mallia%E2%80%99s%20company%20is%20also>

[10] [undercodenews.com](https://undercodenews.com/12-rules-of-agentic-ai-the-hidden-architecture-crisis-behind-enterprise-ai-failure-video/#:~:text=revolution%2C%20yet%20behind%20the%20polished,Meanwhile%2C%20organizations%20rush%20toward) — <https://undercodenews.com/12-rules-of-agentic-ai-the-hidden-architecture-crisis-behind-enterprise-ai-failure-video/#:~:text=revolution%2C%20yet%20behind%20the%20polished,Meanwhile%2C%20organizations%20rush%20toward>

[11] [undercodenews.com](https://undercodenews.com/12-rules-of-agentic-ai-the-hidden-architecture-crisis-behind-enterprise-ai-failure-video/#:~:text=article%20argues%20that%20most%20agentic,created%20by%20John%20Taschek%20of) — <https://undercodenews.com/12-rules-of-agentic-ai-the-hidden-architecture-crisis-behind-enterprise-ai-failure-video/#:~:text=article%20argues%20that%20most%20agentic,created%20by%20John%20Taschek%20of>

Trust and Oversight Take Center Stage

Even as AI agents grow more capable, winning the confidence of users and regulators is essential. A Salesforce survey of 3,200 patients across eight countries finds that people are open to AI help in healthcare, but only on their terms (www.salesforce.com [1]) (www.salesforce.com [2]). Sixty-one percent of global patients say they are comfortable using agent AI in care, and 64% would share their full medical history with AI for a faster diagnosis (www.salesforce.com [3]). However, patients are three times more likely to trust an AI agent in their doctor’s own portal than one from a public chatbot or website (www.salesforce.com [4]) (www.salesforce.com [5]). And nearly 90% insist that any automated assistant offer a simple way to escalate to a human professional when needed (www.salesforce.com [6]). In short, consumers will reward healthcare AI that is convenient and well-supervised, but they reject a black-box approach. Health providers implementing “agent” technology must prioritize transparency, data privacy, and a seamless handoff to human care staff if they want patients to actually use it.

Within companies, a similar principle holds. Many employees — particularly in the US — remain skeptical of workplace AI, flagging low trust, poor training, and unreliable outputs as big sticking points in early pilots (undercodenews.com [7]). These issues have caused a number of corporate AI initiatives to stall. Indeed, industry analyses show that most agent AI failures aren’t caused by the algorithms at all, but by unstable data architectures and lack of attention to governance and change management in deployment (undercodenews.com [8]) (undercodenews.com [9]). The lesson for leaders is that rolling out an autonomous agent is as much about process and people as technology. Employees need to understand how the AI works, see that its recommendations are based on sound data, and trust that the system will seek human input or shut itself off when it hits a risk or uncertainty.

Finally, oversight isn’t just an internal concern — it’s also a growing legal mandate. In a first-of-its-kind ruling, a federal judge in New York has allowed prosecutors to access an executive’s ChatGPT logs as evidence in a fraud case (thelawreporters.com [10]). The defendant tried to block the warrant by claiming his AI conversations were private or attorney–client privileged, but the court found that users “cannot reasonably expect confidentiality” when using a public chatbot service (thelawreporters.com [11]). (Earlier this year, another judge reached a similar conclusion about communications with Anthropic’s Claude AI (thelawreporters.com [12]).) The implication is clear: AI interactions fall under the same compliance and e-discovery rules as any other digital communication. Companies must update policies to prevent sensitive data from leaking via AI tools, and they need to log and monitor agent interactions for potential legal exposure.

All these developments point to the same conclusion: as AI agents become more autonomous and integrated in core workflows, strong governance and user trust are not obstacles — they are

prerequisites for success. It's telling that 77% of tech leaders say AI adoption in their firms is now outpacing their ability to govern it (newsroom.ibm.com [13]). On average, large enterprises dealt with 54 AI "agent incidents" in the past year that required human intervention, and 37% of those cases resulted in data breaches or other security issues (newsroom.ibm.com [14]). Experts advise that unlike traditional IT deployments – where most of the effort is spent upfront – effective AI agent programs need ongoing care after launch, with constant monitoring, validation, and adjustment (www.europesays.com [15]). The payoff for getting this right can be substantial: organizations that bake in rigorous controls and oversight from the start have been found to deploy 16 times more AI agents, see 18% higher operating margins, and achieve far greater ROI on AI initiatives than those taking a move-fast-and-fix-it-later approach (newsroom.ibm.com [16]).

References:

- [1] [www.salesforce.com](https://www.salesforce.com/news/stories/patients-trust-doctors-agents-over-public-ai/#:~:text=Enterprise%20New%20Research%3A%20Patients%20Trust,waited%20on%20hold%20to%20make) — <https://www.salesforce.com/news/stories/patients-trust-doctors-agents-over-public-ai/#:~:text=Enterprise%20New%20Research%3A%20Patients%20Trust,waited%20on%20hold%20to%20make>
- [2] [www.salesforce.com](https://www.salesforce.com/news/stories/patients-trust-doctors-agents-over-public-ai/#:~:text=recommendations,The) — <https://www.salesforce.com/news/stories/patients-trust-doctors-agents-over-public-ai/#:~:text=recommendations,The>
- [3] [www.salesforce.com](https://www.salesforce.com/news/stories/patients-trust-doctors-agents-over-public-ai/#:~:text=recommendations,The) — <https://www.salesforce.com/news/stories/patients-trust-doctors-agents-over-public-ai/#:~:text=recommendations,The>
- [4] [www.salesforce.com](https://www.salesforce.com/news/stories/patients-trust-doctors-agents-over-public-ai/#:~:text=watching%20closely%20to%20see%20if,Patients%E2%80%99) — <https://www.salesforce.com/news/stories/patients-trust-doctors-agents-over-public-ai/#:~:text=watching%20closely%20to%20see%20if,Patients%E2%80%99>
- [5] [www.salesforce.com](https://www.salesforce.com/news/stories/patients-trust-doctors-agents-over-public-ai/#:~:text=matters%20enormously,center%20delivering%20care%20to%20millions) — <https://www.salesforce.com/news/stories/patients-trust-doctors-agents-over-public-ai/#:~:text=matters%20enormously,center%20delivering%20care%20to%20millions>
- [6] [www.salesforce.com](https://www.salesforce.com/news/stories/patients-trust-doctors-agents-over-public-ai/#:~:text=watching%20closely%20to%20see%20if,Patients%E2%80%99) — <https://www.salesforce.com/news/stories/patients-trust-doctors-agents-over-public-ai/#:~:text=watching%20closely%20to%20see%20if,Patients%E2%80%99>
- [7] [undercodenews.com](https://undercodenews.com/12-rules-of-agentic-ai-the-hidden-architecture-crisis-behind-enterprise-ai-failure-video/#:~:text=speed%2C%20or%20model%20capability,and%20readiness%20is%20where%20transformation) — <https://undercodenews.com/12-rules-of-agentic-ai-the-hidden-architecture-crisis-behind-enterprise-ai-failure-video/#:~:text=speed%2C%20or%20model%20capability,and%20readiness%20is%20where%20transformation>
- [8] [undercodenews.com](https://undercodenews.com/12-rules-of-agentic-ai-the-hidden-architecture-crisis-behind-enterprise-ai-failure-video/#:~:text=revolution%2C%20yet%20behind%20the%20polished,Meanwhile%2C%20organizations%20rush%20toward) — <https://undercodenews.com/12-rules-of-agentic-ai-the-hidden-architecture-crisis-behind-enterprise-ai-failure-video/#:~:text=revolution%2C%20yet%20behind%20the%20polished,Meanwhile%2C%20organizations%20rush%20toward>
- [9] [undercodenews.com](https://undercodenews.com/12-rules-of-agentic-ai-the-hidden-architecture-crisis-behind-enterprise-ai-failure-video/#:~:text=article%20argues%20that%20most%20agentic,created%20by%20John%20Taschek%20of) — <https://undercodenews.com/12-rules-of-agentic-ai-the-hidden-architecture-crisis-behind-enterprise-ai-failure-video/#:~:text=article%20argues%20that%20most%20agentic,created%20by%20John%20Taschek%20of>
- [10] [thelawreporters.com](https://thelawreporters.com/chatgpt-attorney-client-privilege-openai-search-warrant#:~:text=A%20US%20federal%20judge%20has,protected%20communications%20with%20his%20lawyers) — <https://thelawreporters.com/chatgpt-attorney-client-privilege-openai-search-warrant#:~:text=A%20US%20federal%20judge%20has,protected%20communications%20with%20his%20lawyers>
- [11] [thelawreporters.com](https://thelawreporters.com/chatgpt-attorney-client-privilege-openai-search-warrant#:~:text=prosecutorial%20scrutiny) — <https://thelawreporters.com/chatgpt-attorney-client-privilege-openai-search-warrant#:~:text=prosecutorial%20scrutiny>
- [12] [thelawreporters.com](https://thelawreporters.com/chatgpt-attorney-client-privilege-openai-search-warrant#:~:text=The%20case%20parallels%20a%20February,not%20covered%20by%20attorney%E2%80%93client%20privilege) — <https://thelawreporters.com/chatgpt-attorney-client-privilege-openai-search-warrant#:~:text=The%20case%20parallels%20a%20February,not%20covered%20by%20attorney%E2%80%93client%20privilege>
- [13] [newsroom.ibm.com](https://newsroom.ibm.com/2026-06-08-new-ibm-study-finds-cios-and-ctos-face-growing-ai-control-gap-as-enterprise-deployment-scales#:~:text=CEO,redesigning%20how%20organizations%20control%2C%20govern) — <https://newsroom.ibm.com/2026-06-08-new-ibm-study-finds-cios-and-ctos-face-growing-ai-control-gap-as-enterprise-deployment-scales#:~:text=CEO,redesigning%20how%20organizations%20control%2C%20govern>
- [14] [newsroom.ibm.com](https://newsroom.ibm.com/2026-06-08-new-ibm-study-finds-cios-and-ctos-face-growing-ai-control-gap-as-enterprise-deployment-scales#:~:text=and%20invest%20in%20it%20and,of%20those%20AI%20agent%20incidents) — <https://newsroom.ibm.com/2026-06-08-new-ibm-study-finds-cios-and-ctos-face-growing-ai-control-gap-as-enterprise-deployment-scales#:~:text=and%20invest%20in%20it%20and,of%20those%20AI%20agent%20incidents>
- [15] [www.europesays.com](https://www.europesays.com/ai/85143/#:~:text=deployments%2C%20Salesforce%20has%20identified%20many,incloding%20managing%20and%20improving%20them) — <https://www.europesays.com/ai/85143/#:~:text=deployments%2C%20Salesforce%20has%20identified%20many,incloding%20managing%20and%20improving%20them>
- [16] [newsroom.ibm.com](https://newsroom.ibm.com/2026-06-08-new-ibm-study-finds-cios-and-ctos-face-growing-ai-control-gap-as-enterprise-deployment-scales#:~:text=years%2C%20raising%20the%20stakes%20for,portable%20and%20models%20replaceable%20rather) — <https://newsroom.ibm.com/2026-06-08-new-ibm-study-finds-cios-and-ctos-face-growing-ai-control-gap-as-enterprise-deployment-scales#:~:text=years%2C%20raising%20the%20stakes%20for,portable%20and%20models%20replaceable%20rather>

Key Statistics

- Anthropic says 65% of its product team's code is now generated by its internal AI agent ([venturebeat.com])(<https://venturebeat.com/technology/anthropic-launches-claude-tag-replacing-its-slack-app-with-a-persistent-ai-teammate-that-learns-monitors-and-works-autonomously#:~:text=implications%20for%20enterprise%20workflow%2C%20governance%2C,How%20Claude%20Tag%20works>)).
- Patients are 3x more likely to trust an AI agent in their doctor's portal than one on a public site ([www.salesforce.com])(<https://www.salesforce.com/news/stories/patients-trust-doctors-agents-over-public-ai/#:~:text=watching%20closely%20to%20see%20if,Patients%E2%80%99>) ([www.salesforce.com])(<https://www.salesforce.com/news/stories/patients-trust-doctors-agents-over-public-ai/#:~:text=matters%20enormously,center%20delivering%20care%20to%20millions>)).
- 89% of patients insist on an 'escalate to human' option before they'll trust an AI for support tasks ([www.salesforce.com])(<https://www.salesforce.com/news/stories/patients-trust-doctors-agents-over-public-ai/#:~:text=2026%205%20min%20read%20Key,Patients%E2%80%99>)).

- Tech executives anticipate a 38% increase in the number of AI agents deployed in their enterprises by 2027 ([newsroom.ibm.com](https://newsroom.ibm.com/2026-06-08-new-ibm-study-finds-cios-and-ctos-face-growing-ai-control-gap-as-enterprise-deployment-scales#:~:text=business%20are%20deploying%20technology%20faster,Governance%20is%20also)).
- Organizations reported an average of 54 AI-related “agent incidents” in the past year, with 37% involving data exposure or security breaches ([newsroom.ibm.com](https://newsroom.ibm.com/2026-06-08-new-ibm-study-finds-cios-and-ctos-face-growing-ai-control-gap-as-enterprise-deployment-scales#:~:text=and%20invest%20in%20it%20and,of%20those%20AI%20agent%20incidents)).

KEY TAKEAWAY

AI agents are rapidly advancing from simple Q&A bots to true autonomous assistants that can execute business operations. Leaders should move from pilots to governed deployments, adding oversight and domain-specific depth to harness these tools safely and at scale.

Sources

[Anthropic launches Claude Tag, replacing its Slack app with a persistent AI teammate that learns, monitors and works autonomously](https://venturebeat.com/technology/anthropic-launches-claude-tag-replacing-its-slack-app-with-a-persistent-ai-teammate-that-learns-monitors-and-works-autonomously)

<https://venturebeat.com/technology/anthropic-launches-claude-tag-replacing-its-slack-app-with-a-persistent-ai-teammate-that-learns-monitors-and-works-autonomously>

[Backbase Acquires Kasisto To Advance Agentic Banking](https://pulse2.com/backbase-acquires-kasisto-to-advance-agentic-banking/)

<https://pulse2.com/backbase-acquires-kasisto-to-advance-agentic-banking/>

[Syndio bets on agentic AI with first acquisition in Seattle pay equity startup's history](https://www.geekwire.com/2026/syndio-bets-on-agentic-ai-with-first-acquisition-in-seattle-pay-equity-startups-history/)

<https://www.geekwire.com/2026/syndio-bets-on-agentic-ai-with-first-acquisition-in-seattle-pay-equity-startups-history/>

[New Research: Patients Trust Their Doctor's AI Agents 3x More Than Public AI](https://www.salesforce.com/news/stories/patients-trust-doctors-agents-over-public-ai/)

<https://www.salesforce.com/news/stories/patients-trust-doctors-agents-over-public-ai/>

[US Judge Allows Search Warrant Targeting Executive's AI Chatbot Communications](https://thelawreporters.com/chatgpt-attorney-client-privilege-openai-search-warrant)

<https://thelawreporters.com/chatgpt-attorney-client-privilege-openai-search-warrant>

[New IBM Study Finds CIOs and CTOs Face Growing AI Control Gap as Enterprise Deployment Scales](https://newsroom.ibm.com/2026-06-08-new-ibm-study-finds-cios-and-ctos-face-growing-ai-control-gap-as-enterprise-deployment-scales)

<https://newsroom.ibm.com/2026-06-08-new-ibm-study-finds-cios-and-ctos-face-growing-ai-control-gap-as-enterprise-deployment-scales>

[Exclusive: Seltz, a startup rebuilding web search for AI agents, raises \\$12.5 million in seed funding](https://finance.yahoo.com/technology/ai/articles/exclusive-seltz-startup-rebuilding-search-110000196.html)

<https://finance.yahoo.com/technology/ai/articles/exclusive-seltz-startup-rebuilding-search-110000196.html>

[Agentic infrastructure startup Seltz raises \\$12.5M to help AI agents search the web for answers](https://siliconangle.com/2026/06/24/agentic-infrastructure-startup-seltz-raises-12-5m-help-ai-agents-search-web-answers/)

<https://siliconangle.com/2026/06/24/agentic-infrastructure-startup-seltz-raises-12-5m-help-ai-agents-search-web-answers/>

[12 Rules of Agentic AI: The Hidden Architecture Crisis Behind Enterprise AI Failure](https://undercodenews.com/12-rules-of-agentic-ai-the-hidden-architecture-crisis-behind-enterprise-ai-failure-video/)

<https://undercodenews.com/12-rules-of-agentic-ai-the-hidden-architecture-crisis-behind-enterprise-ai-failure-video/>

