

# Predicts 2026: AI Agents Will Transform IT Infrastructure and Operations

4 December 2025 - ID G00836214 - 23 min read

By: Cameron Haight, Ashish Banerjee, Joe Antelmi, Jonathan Forest, Chandra Mukhyala

Initiatives: [I&O Organizational Strategy](#)

AI agents are beginning to be deployed in today's I&O organizations and their usage will only increase in the next few years. To prepare, it is crucial for heads of I&O to assess not only their benefits, but to also remain aware of their continuing challenges to ensure safe and successful production implementations.

## Overview

### Key Findings

- Heads of I&O will face pressure to replace human staff with AI-agent staff to enable agent system investments with less overall impact on their budgets, leading to the increased commonality of diamond-shaped organizations where AI agents replace human administrator engineers. This is also driven by the increasing capability of advanced AI services that can perform complex tasks exceeding human expertise.
- AI is evolving from assistants to more autonomous agentic systems capable of planning, decision-making and goal fulfillment. This will reshape the IT operations landscape by moving from tools that assist humans to platforms that replace manual effort for complex workflows.
- Agentic NetOps is projected to be more scalable and cost-effective than current manual operations and outsourcing. It can offer transformational benefits not realized in previous AI and automation technologies by minimizing humans required to be “in the loop.” This capability is expected to lead to a shift of enterprises from MNS to DIY management
- AI agents are rapidly assuming mission-critical roles, which necessitates formal performance management frameworks to mitigate operational and compliance risks. The allocation of operational risk and accountability to technology providers will fundamentally alter vendor management practices and contractual relationships.
- Heads of I&O, frustrated by the limits of traditional automation tools, will increasingly look to agentic AI to automate complex, nondeterministic IT infrastructure operations with a human language interface. This approach simplifies adoption by providing natural language syntax, reducing the reliance on specialized DevOps skills.

## Recommendations

- Prepare humans to change roles, as many manual administration and complex migration/rehosting automation tasks could be replaced by AI agents. Additionally, invest in an AI agent engineering platform to manage the inherent risks of AI agent adoption.
- Rethink talent and skills development. The need for manual operators will diminish, while demand for AI strategists, trainers, engineers and governance specialists will rise. Heads of I&O should invest in upskilling their teams, focusing on data engineering, AI model management, and strategic thinking, transforming the role of the IT operator from a “doer” to a “supervisor” of AI agents.
- Calculate the business case by determining the total cost of ownership (TCO) of investing in agentic NetOps in a DIY model versus outsourcing to MNS providers. Furthermore, start reallocating network spending by shifting some network budgets away from managed network services (MNS) to agentic NetOps software licenses and products.
- Update vendor contracts to include enforceable AI agent performance SLAs with clear remediation, audit, escalation procedures, and financial remedies. This includes collaborating with vendors to define standardized, AI-specific SLA metrics.
- Transform infrastructure administrator roles into agentic AI automation engineer and supervisor role, leveraging their deep domain experience, knowledge, and adaptability to introduce and mature agentic AI-driven operations. This transformation will involve shifting I&O’s focus from performing tasks to teaching, supervising, and setting guardrails as well as success goals for AI agents.

## Strategic Planning Assumptions

By 2030, 50% of I&O organizations will be reshaped as heads of I&O invest in AI agents to accomplish complex I&O management tasks.

By 2028, human-in-the-loop of agentic AI workflows in IT operations will decrease to 40%, down from 95% in 2025.

By 2030, 20% of organizations will shift from MNS to DIY after adopting agentic NetOps, resulting in lower costs by at least 20%.

By 2028, 50% of I&O leaders will require AI agent performance SLAs in vendor contracts, up from less than 5% in 2025.

By 2029, 70% of enterprises will deploy agentic AI as a part of IT infrastructure operations, compared to less than 5% in 2025.

## Analysis

### What You Need to Know

The enterprise landscape is on the cusp of a significant transformation driven by the rapid adoption and evolution of agentic AI as CIOs and technology leaders are prioritizing investment in AI agents. A recent Gartner survey identified AI agents as the top technology that CIOs and technology leaders expect to increase their investment in for the remainder of 2025. <sup>1</sup> However, organizations are not just planning for AI agent adoption, they are actively implementing it. A majority of organizations (60%) are already either piloting or scaling use cases for AI agent automation embedded within existing software. <sup>2</sup>

Agentic AI is expected to fundamentally alter how organizations operate. A significant 66% of surveyed participants believe that agentic AI, with its capability to autonomously plan and act, will change how their organizations operate by 2028. <sup>3</sup> This goes beyond mere automation potentially pointing to a shift in core operational paradigms. Consequently, nearly 70% of organizations will start managing an AI agent workforce within the next 5 years highlighting the emergence of AI agents as a distinct “workforce” requiring management and oversight. <sup>4</sup>

From our client discussions, heads of I&O are beginning to assess how to embed AI agents that can coordinate tasks ranging from customer support to network operations. At the same time, technology providers are rushing to host embedded agents within their platforms such as in the case of ITSM suites, observability and monitoring tools as well as I&O automation and orchestration technologies. While most of these implementations are addressing relatively simple use cases, vendors are also experimenting with multiagent designs that enable more complex processing including cross-functional workflow orchestration.

While optimism is high regarding their potential to streamline workflows and enhance productivity, organizations face considerable challenges related to issues such as integration with legacy systems, ensuring correct data for processing, security and compliance as well as managing their dynamic cost impacts. Robust governance frameworks and aggressive skills development will be two of the critical areas for heads of I&O to address in order to effectively navigate this new AI-driven paradigm successfully.


## Strategic Planning Assumptions

**Strategic Planning Assumption:** By 2030, 50% of I&O organizations will be reshaped as heads of I&O invest in AI agents to accomplish complex I&O management tasks.

**Analysis by:** Joe Antelmi

### Key Findings:

- Services built upon complex, specialized AI agent architectures, such as Claude Code, are increasing in capability. The best LLM-based AI agent systems are not just assistants, they are teachers and mentors.
- Trusted, sophisticated AI agent systems built up on AI agent platforms will be able to accomplish complex tasks more quickly than humans, because they will be able to massively iterate parallelized agents and select the highest quality output. These complex use cases could include such as refactoring infrastructure into cloud native services or handling policy as code conflicts for a complex on-premises environment.
- Successful I&O implementers of AI agents, built or bought, will need to invest in sophisticated AI agent platform engineering services. For bought AI agents, these platform capabilities will include managing AI agent permissions, integrations, testing, evaluation, guardrails, as well as nonfunctional requirements such as backup, scaling, failover, and cost management. For internally developed AI agents, these platforms will also need to support internal stakeholders with provisioning, automating and scaling infrastructure to support AI agent services.
- I&O organizations are heavily dependent on their vendor partners for AI capabilities, as the skills and expertise to scale AI agent systems more likely exist inside cloud, AI and I&O software and hardware vendors than inside the I&O organization. Heads of I&O will need to find room in mostly flat budgets to invest in AI agents, and some of that budget may come from replacing human salaries with AI-agent licenses. Diamond-shaped organizations thus will become more common (see Figure 1) where AI agents replace human operators, and engineers.

**Figure 1: Future Reshaped I&O Organization With AI Agents****Future Reshaped I&O Organization with AI Agents**

Source:  
836214

**Gartner.**

**Near-Term Flag:**

In the near term, expect to observe I&O professionals using advanced AI services accomplishing complex tasks in minutes that would take them hours. Expect all of your software vendor partners to try to sell you an AI platform and specific AI agent services. Expect to have trouble showing ROI for AI investments due to the upfront cost of these investments, vs modest productivity gains without more extreme restructuring or transformation. Expect certain parts of your budget to receive scrutiny, such as outsourcing and service integrators, due to the perception that these domains can largely be replaced by AI services (whether or not this is actually true).

**Market Implications:**

Heads of I&O will have a mix of vendor partners and relationships, some will enable, and some will constrain the AI agent future in their organization. In the best case scenarios, I&O and AI and Cloud vendors are partners for heads of I&O in this journey. In other common scenarios, aggressive software vendors will acquire existing licensed products and bundle, audit, and increase their prices in a way that strangles the budget for AI innovation, unless it is linked to increasing investments in that software vendor.

I&O software vendors will increasingly change their licensing models to make it harder for internal I&O organizations to recreate their software using vibe coding, while providing heavy license incentives to adopt AI agent platforms and services so that they are closely tied to the value and the implementation of AI agents.

## Recommendations:

- Plan now to ensure your future contract renewals are with vendors that will support, rather than inhibit, the AI agent future of I&O Vendors with flexible and transparent pricing models that enable usage-based or outcome-based pricing should be prioritized.
- Look to hire software and AI engineers into I&O roles to become the future AI-enabled engineers and AI agent orchestrators. Moreover, hire these roles into a SRE style niche, where they focus on enabling the platforms that enable I&O for the future.
- Invest in an AI agent engineering platform or else beware of the extreme risks that AI agent adoption provides to security, identity management and cost overruns. If you don't build and sustain an AI platform engineering function in collaboration with other teams, you will be left out.
- Prepare humans to change roles, as many of their manual administration and even complex migration rehosting automation tasks could be replaced by AI agents.

## Related Research:

[How to Securely Delegate Access From Humans to AI Agents](#)

[The Future of I&O Automation Is Agentic So Begin Piloting Now](#)

[How to Structure Your Infrastructure Platform Engineering Journey](#)

## Top Strategic Technology Trends for 2026: AI-Native Development Platforms

**Strategic Planning Assumption:** By 2028, human-in-the-loop of agentic AI workflows in IT operations will decrease to 40%, down from 95% in 2025.

**Analysis by:** Ashish Banerjee

### Key Findings:

- At present, AI agents automate simple, well-defined tasks with low, conditional autonomy and basic reasoning for multistep planning. They rely on reactive, retrieval-augmented perception, flexible but limited plug-and-play APIs, and primarily offline RLHF learning with short-term, session-only memory.
- AI is transitioning from early pretrained models that provided basic text generation and few-shot responses to more reasoning-based systems capable of planning, decision-making and goal fulfillment.
- Future AI agents will open new use cases by offering greater autonomy, handling complex workflows, providing domain expertise and multimodal understanding, and using advanced tools for execution and collaboration. Consequently, the role of the IT operator will transform from a “doer” to a “supervisor” of AI agents. The decline in human-in-the-loop will initially result in role elimination, but ultimately lead to role evolution fueled by rehiring.

### Market Implications:

The transition from human-supervised, reactive AI agents to highly autonomous systems capable of independent planning and execution will reshape the IT operations landscape. This shift will move the market from tools that assist humans to platforms that replace manual effort for complex workflows.

- **Evolution of AIOps platforms:** The current generation of AIOps tools, which focus on data aggregation, anomaly detection, incident management and root cause analysis, will evolve into autonomous operations platforms. These future systems will not just identify problems but will independently execute complex resolution workflows across multiple domains, from networking to cloud infrastructure management. As organizations pivot toward these capabilities, I&O budgets will likely shift in favor of AIOps and agent platforms, responding to anticipated strong market growth. In this evolving landscape, platform selection and integration strategy will emerge as key competitive differentiators, potentially reshaping how enterprises achieve operational excellence and maintain a strategic edge.
- **Consolidation of monitoring and automation tools:** As AI agents become capable of using tools and collaborating, the need for a multitude of niche monitoring and automation solutions will decrease. Vendors will likely consolidate their offerings into integrated platforms where the AI agent serves as the central orchestrator, reducing tool sprawl and simplifying the IT environment. Success of IT operations would hinge on how well AI agent platforms are able to onboard unified telemetry/observability, event correlation, and orchestration hooks across CI/CD, ITSM, and infra APIs.
- **Emergence of new service models:** Managed service providers (MSPs) and vendors will introduce new “autonomous IT operations” services. These offerings will provide pretrained AI agents with domain-specific expertise, capable of managing entire segments of the IT stack with minimal human oversight. This will create new opportunities for outcome-based pricing models, where payment is tied to performance such as measurable improvements in MTTR, uptime, automation coverage, and SLO attainment rather than human hours, licenses or generic feature claims.
- **Increased focus on proactive and predictive operations:** With AI agents handling complex, multistep tasks, the focus of IT operations will shift from reactive troubleshooting to proactive and predictive maintenance. AI will anticipate potential issues and take corrective action before they impact business services, leading to gains in reliability and resilience. Runbooks will evolve into self-healing workflows as agents take on Tier 1 and Tier 2 tasks like alert noise suppression, triage, correlation, and standardized remediation, with humans stepping in for edge cases and control points.

## Recommendations:

- Start with low-risk, well-defined, high-volume and high-impact tasks that can be fully automated, such as alert deduplication, incident classification, known-error remediation and scaling actions. Use incident data to build a prioritized backlog of quick wins with clear success metrics and guardrails. As your team gains confidence, gradually expand to more complex, cross-domain workflows.
- Revisit your AI strategy. While advanced agent capabilities such as full autonomy, ecosystemwide actions and collective behavior represent long-term goals, current capabilities remain limited. Focus on monitoring developments in these areas and identifying incremental opportunities to enhance your existing systems, laying the groundwork for future adoption as enabling technologies mature.
- Rethink talent and skills development. The need for manual operators will diminish, while demand for AI strategists, trainers, engineers, “red teamers” and governance specialists will rise. Heads of I&O should invest in upskilling their teams, focusing on elementary data science, AI model management, and strategic thinking. Prioritize reskilling into SRE, AIOps engineering, automation ownership, policy-as-code administration, and model risk management. Pair deep domain expertise with agent orchestration skills.
- Prioritize governance and trust. Establish an agent governance framework covering policy-as-code, approval ladders, observability of agent actions, and audit logging mapped to change controls. Develop clear policies for AI decision making, risk management, and oversight. Building trust in these autonomous systems will be essential for adoption, requiring transparency in how AI agents operate and make decisions. Implement guardrails, auditability, rollback plans, and value-led prioritization to avoid hype traps.
- Engage with vendors as strategic partners. Evaluate vendors not just on their current capabilities but on their long-term vision for autonomous operations, and their ability to execute. Partner with those who are investing in the core technologies you’ve identified — such as advanced planning, multimodal understanding, and collaborative AI — to ensure your technology stack remains future-proof. Demand proofs tied to production-like sandboxes, outcome SLAs, and clear safety interlocks; scrutinize claims for genuine planning/execution autonomy versus rebranded assistants or RPA.

## Related Research:

[Reengineer I&O Processes With Integrative Agentic AI](#)

## Emerging Tech: AI Vendor Race: The Key Defining Characteristics of Agentic AI

### Emerging Tech: AI Vendor Race: Agentic AI Maturity Roadmap

**Strategic Planning Assumption:** By 2030, 20% of organizations will shift from MNS to DIY after adopting agentic NetOps.

**Analysis by:** Jonathan Forest

#### Key Findings:

- The use of AI agents for network operations (that is, agentic NetOps) is a hyped concept, but vendor capabilities are either nonexistent or still limited to Day 2 use cases with “humans in the loop.” Vendors are investing heavily in this space, as we expect new offerings to be introduced to the market over the coming months.
- Up until now, AI and automation have had limited adoption in networking due to complexity, unreliability of outputs and unclear business value. In fact, Gartner estimates that only one-third of network tasks are automated, and AI networking is used to a lesser extent.
- Many enterprises rely on managed network service providers (MNS) to outsource network management due to internal skills and resource gaps. Addressing skills, resource limitations and optimizing costs are go forward challenges for network operations teams.
- Agentic NetOps can offer transformational benefits not realized in previous AI and automation technologies by minimizing humans “from the loop” resulting in a shift of enterprises from MNS to do-it-yourself (DIY) management.

#### Near-Term Flag:

We expect significant functionality to be delivered by vendors in the next 12 months. This will drive more piloting of solutions by end user organizations that will lead to adoption. Additionally, some functionality will exist as part of network vendor offerings which will inherently further drive adoption.

#### Market Implications:

Agentic NetOps leverages goal-driven autonomous AI agents that have been granted rights by the organization to operate network tasks and processes independently. AI agents may work independently or as part of a system to manage network infrastructure life cycle management with minimal to no human involvement.

This is different from previous versions of AI as agentic NetOps can operate independently from and on behalf of humans. Simply put, software acts as an agent to perform networking tasks that are programmed to operate independent of “human in the loop” involvement to more “human on the loop” acting in a more supervisory and validation capacity. Agentic NetOps is meant to replace or augment network operations roles/processes/functions to drive more efficiency, speed, agility and accuracy of outcomes. AI agents may be hosted on or extract data from any networking device/component (e.g., switches, routers, SD-WAN devices, access points and firewalls) and tools and may have different and specialized roles. For example, different AI agents may be responsible for monitoring traffic, enforcing security policies or for troubleshooting. These AI agents interact with large language models (LLMs) or domain specific small language models (SLMs) to gain the knowledge/insights for the actions to take. They can accumulate knowledge/memories as context, over time, which influences their behavior and decision making based on more inputs.

Agentic NetOps are and will be delivered by the following types of vendors/providers:

1. Integrated as part of network/security vendor offerings (e.g., routers, switches, APs, Firewalls, SD-WAN, SSE, SASE).
2. Multivendor over the top solutions that integrate with network/security vendor offerings.
3. MNS providers that offer an agentic NetOps service.

As this functionality becomes more available in existing networking products and tooling, the option to manage networks internally in a do-it-yourself (DIY) mode becomes more feasible and reduces the needs for MNS offerings. Agentic NetOps is expected to be more scalable and cost-effective than current manual operations and outsourcing to MNS providers. In fact, Gartner expects that organizations who fail to adopt agentic NetOps will overspend by 25% on network management by 2030.

Ultimately, this is the turning point of AI moving from being a human support tool to operating processes independently but still directed by humans with guardrails, access rights and goals/objectives defined. We expect this will lead to a shift in enterprise network budgets/spend to more software that delivers agentic NetOps and away from managed network services (MNS) and support. This will take some time as vendor agentic NetOps capabilities improve to provide similar capabilities as MNS providers. During this time, we expect enterprises to gradually add more agentic NetOps capabilities to address certain tasks/processes and functions and still retain some MNS capabilities. But, once the agentic NetOps capability is fully rolled out and adopted by enterprises, we expect more organizations to operate in a DIY mode.

## Recommendations:

- Test and validate agentic NetOps functionality by confirming outcomes, such as root cause analysis and recommended actions, via pilots, and speak with reference customers. Allocate innovation time for personnel to investigate and pilot agentic NetOps solutions.
- Calculate the business case by determining the TCO of investing in agentic NetOps in a DIY model vs. outsourcing to MNS providers as well as comparing functionality.
- Address skills and resource gaps in network operations by strategically adopting agentic NetOps functionality.
- Start reallocating network spending by shifting some network budgets away from managed network services (MNS) to agentic NetOps software licenses and products.

## Related Research:

[Why Agentic NetOps Will Move Enterprises Away From Managed Network Services](#)

[Agentic NetOps Will Revolutionize Network Operations](#)

**Strategic Planning Assumption:** By 2028, 50% of I&O leaders will require AI agent performance SLAs in vendor contracts, up from less than 5% in 2025.

**Analysis by:** Cameron Haight

## Key Findings:

- AI agents are rapidly assuming mission-critical roles, necessitating formal performance management frameworks to mitigate operational and compliance risks.
- There is a significant gap between the current state of AI agent SLA adoption and the projected requirement driven by both technological and contractual immaturity.
- The allocation of operational risk and accountability to technology providers will fundamentally alter vendor management practices and contractual relationships.

## Market Implications:

The formalization of AI agent performance SLAs in enterprise vendor contracts will likely catalyze significant shifts across the technology provider landscape potentially altering the historical processes of procurement and ongoing service management. As AI becomes embedded in mission-critical operations, enterprises will demand greater transparency, reliability and accountability from vendors. If this evolves, technology providers will be compelled to invest in advanced operational monitoring, reporting and explainability features to meet increasingly stringent SLA requirements. This could result in new innovation focused on AI observability platforms, automated performance analytics and real-time incident response capabilities. Vendors unable to demonstrate robust SLA compliance could face diminished competitiveness particularly in regulated industries such as financial services and healthcare.

The proliferation of AI agent SLAs may contribute to the development of industrywide standards and certification programs as enterprises seek consistency and comparability in vendor offerings. Certification will become a market differentiator enabling vendors to demonstrate their commitment to accountability and best practices and providing enterprises with greater confidence in vendor selection. Procurement processes could evolve to prioritize vendors with demonstrable SLA compliance and operational transparency. RFPs and vendor evaluations will increasingly include requirements for AI agent performance metrics, historical SLA attainment as well as possibly detailed remediation capabilities. Contract negotiations will likely become more complex with enterprises demanding granular SLA definitions, audit rights and financial penalties for noncompliance. This will necessitate closer collaboration between IT, procurement and legal functions within the enterprise.

## Recommendations:

- Collaborate with vendors and (and potentially) industry organizations to define standardized, AI-specific SLA metrics encompassing accuracy, timeliness, explainability and resilience.
- Institute rigorous vendor evaluation processes that assess the maturity of AI operational monitoring, reporting and compliance capabilities.
- Update vendor contracts to include enforceable AI agent performance SLAs with clear remediation, audit escalation procedures, and financial remedies.

## Related Research:

- [The Current State of AI Agents for Enterprises](#)
- [Tool: AI Agent Assessment Framework](#)

**Strategic Planning Assumption:** By 2029, 70% of enterprises will deploy agentic AI to semiautonomously operate their IT infrastructure, compared to less than 5% in 2025

**Analysis by:** Chandra Mukhyala

## Key Findings:

- Frustrated by the limits of traditional automation tools, heads of I&O will increasingly look to agentic AI to automate complex, nondeterministic IT infrastructure operations with a human language interface.
- Unlike traditional AI automation in I&O, which heavily relies on APIs and requires specialized DevOps skills that are often costly and scarce, agentic AI simplifies adoption by using a natural language syntax.
- Heads of I&O see the promise of agentic AI for automating infrastructure operations, but their skepticism about the cost, complexity and risk of implementation is preventing them from taking action.

**Market Implications:** First, I&O staff should be aware that this trend will fundamentally transform the role of I&O from a manual operator into an AI orchestrator and supervisor. The assumption is not that AI will replace I&O staff, but that it will automate the repetitive, high-volume tasks that currently consume the time of I&O, such as initial incident diagnostics, running routine health checks, and executing standard remediation playbooks. The value of I&O will shift dramatically from performing the task to teaching, supervising, and setting the guardrails for the AI agents that do the work. The deep domain expertise within I&O in networking, storage, or computing becomes more critical than ever, as it will be the knowledge of I&O staff that is used to train, refine, and safely govern these semiautonomous systems.

Second, the tools and interfaces I&O uses daily will evolve significantly. The vendors serving I&O will increasingly embed AI directly into I&O management platforms. Initially, this will take the form of “AI Assistants” that I&O can query in natural language to perform complex diagnostics in seconds. As the technology matures, I&O will be configuring and managing vendor-provided “Specialist Agents” to handle tasks like automated data tiering or security patching within policies defined by I&O. Ultimately, I&O will be using “Agent Development Kits” to build custom agents, effectively codifying the operational expertise of I&O to solve unique challenges within the I&O environment. The I&O interaction model will shift from CLI commands and scripts to a combination of prompt engineering, policy definition, and workflow orchestration.

Finally, this shift requires I&O to proactively invest in a new set of skills to remain relevant and valuable. I&O must move beyond managing infrastructure to managing the AI that manages the infrastructure. The focus should be on developing competencies in agent orchestration, AI security fundamentals (like defending against prompt injection), and understanding how to monitor the performance and cost of the AI systems themselves. The most critical and highest-value role for I&O staff will be that of the expert exception handler — the person who solves the novel, complex problems that fall outside the AI’s training data. I&O should start now by experimenting with the AI and automation features within existing I&O tools and championing a culture of safe, sandboxed experimentation on the I&O team.

## **Recommendations:**

- Transform infrastructure administrator roles into an agentic AI automation and supervisor engineer role by using their domain experience, knowledge and adaptability to introduce and mature agentic AI-driven operations.

- Demand early access to vendor AI features by actively engaging with your infrastructure vendors to test and provide feedback on their new agentic tools.
- Shift from reactive fixing to proactive automation by codifying your operational knowledge into automated playbooks that new AI agents can safely execute.
- Embrace a governance role for AI agents by defining and implementing the policy-based guardrails that ensure all automated actions are safe and auditable.
- Use automated enrichment to economically cleanse, cluster, and tag mass telemetry; reduce noise and expand AI visibility.
- Manually curate high-fidelity telemetry for critical business services to ensure trusted, low-latency AI agent decision-making.

## Related Research:

[Use Agentic AI to Deliver Semiautonomous IT Infrastructure Operations](#)

## A Look Back

*In response to your requests, we are taking a look back at some key predictions from previous years. We have intentionally selected predictions from opposite ends of the scale – one where we were wholly or largely on target, as well as one we missed.*

This report is too new to have on-target or missed predictions.

## Acronym Key and Glossary Terms

|      |                                                                                                                                                                                                                                                                                                                                                                                                      |
|------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| RLHF | Reinforcement learning from human feedback is a machine learning technique that combines automated and human guidance to train an AI agent or a model. It includes training a reward model, based on positive and negative human feedback, then using this model to optimize the AI agent to better understand natural language, learn from human experience and better align with human preference. |
|------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

## Evidence

<sup>1</sup> 2025 Gartner AI Survey: CIO and Technology Leader View

<sup>2</sup> 2025 Gartner Cybersecurity Innovations in AI Risk Management and Use Survey

<sup>3</sup> Gartner Peer Community Poll

## Recommended by the Authors

Some documents may not be available as part of your current Gartner subscription.

[Reengineer I&O Processes With Integrative Agentic AI](#)

[The Future of I&O Automation Is Agentic So Begin Piloting Now](#)

[Why Agentic NetOps Will Move Enterprises Away From Managed Network Services](#)

[Use Agentic AI to Deliver Semiautonomous IT Infrastructure Operations](#)

---

© 2025 Gartner, Inc. and/or its affiliates. All rights reserved. Gartner is a registered trademark of Gartner, Inc. and its affiliates. This publication may not be reproduced or distributed in any form without Gartner's prior written permission. It consists of the opinions of Gartner's Business and Technology Insights Organization, which should not be construed as statements of fact. While the information contained in this publication has been obtained from sources believed to be reliable, Gartner disclaims all warranties as to the accuracy, completeness or adequacy of such information. Although Gartner insights may address legal and financial issues, Gartner does not provide legal or investment advice and its insights should not be construed or used as such. Your access and use of this publication are governed by [Gartner's Usage Policy](#). Gartner prides itself on its reputation for independence and objectivity. Its insights is produced independently by its Business and Technology Insights Organization without input or influence from any third party. For further information, see "[Guiding Principles on Independence and Objectivity](#)." Gartner insights may not be used as input into or for the training or development of generative artificial intelligence, machine learning, algorithms, software, or related technologies.