

EBOOK

AI Challenges Threatening IT Leader Success

Insights from 1,001 enterprise IT leaders
on the new challenges of AI governance

Table of Contents

Executive summary	2
CHALLENGE 1 "Our governance can't keep pace with adoption."	3
CHALLENGE 2 "I don't know what's happening in my own organization."	4
CHALLENGE 3 "I'm being held accountable for systems I can't control."	5
CHALLENGE 4 "Our CEO and board are misaligned, and external pressure is making it worse."	6
CHALLENGE 5 "I'm desperate for observability and auditability."	7
What this means	8
Methodology	9

Executive summary

AI adoption is accelerating across business operations. Last year, the fear was being left behind in the AI revolution; now, the worry is that adoption is outpacing governance. The looming threat is that IT leaders will be held accountable if, and when, something goes wrong.

To better understand today's landscape, Workato partnered with Censuswide to survey 1,001 US IT decision-makers, all VP-level and above in IT, Technology, Engineering, Data, or AI functions, at organizations with over 1,000 employees.

The results uncovered significant concerns about the risk, observability, and ultimate accountability associated with AI technology. In fact, three-quarters of the IT leaders interviewed fear that it would be them, not the technology vendor, who would be held responsible should an AI-related incident occur, potentially damaging their present and future career prospects.

Yet this accountability dynamic creates a troubling tension. When asked what would happen if an AI agent caused a major security or compliance incident tomorrow, more than one in three IT leaders believe their organization would "attempt to handle it quietly without disclosure."

IT leaders are expected to govern systems they don't control, prevent incidents they can't foresee, and ultimately absorb the career risk when their organizations prioritize damage control over transparency.

95%

believe that with the right governance in place, the benefits of agentic AI and digital coworkers are overwhelmingly positive.

Despite this fear, and in contrast with the anti-AI news cycles, nearly all of those questioned, over 95%, believe that with the right governance in place, the benefits of agentic AI and digital coworkers are overwhelmingly positive. However, they do feel that some things need to change to be successful. This report covers the five biggest confessions uncovered in our research.

CHALLENGE 1

Our governance can't keep pace with adoption.

For many of the IT leaders surveyed, the speed of AI adoption is outpacing their ability to govern effectively, which is significantly impacting their overall confidence and trust in the technology.

AI has benefits to business growth, but also some risks. To maximize benefits while minimizing risks, IT leaders know they must establish control and execution standards for AI across the organization.

IT leaders are investing in governance and budgeting in parallel

IT leaders have made a serious commitment to AI governance. A massive 92% of IT leaders expect their organizations to adopt digital coworkers over the next 18 months. Eighty-seven percent have already implemented or are actively developing governance frameworks and have allocated budget specifically for AI agent governance, monitoring, or security tooling.

Yet despite this investment, IT is losing ground on AI agent control

At the same time, a majority of IT leaders (64%) expect unauthorized deployments to accelerate alongside this trend, thereby extending the risk surface.


When asked where AI agents are entering the organization, we learned that adoption is currently coming from multiple sources. Sixty-six percent of agent adoption is coming through IT's formal procurement process. The other 33% are deployed directly, without oversight, by business units and individual employees.

Why? Policies that can't be enforced create a veneer of control while adoption is driven underground into the shadows of IT. Centralized approval creates friction, and as history has shown us, innovation always finds a way to move around friction.

Eager end-users in Sales, HR, IT, and other departments are learning to use AI agents, approved and unapproved, through hands-on experience in their day-to-day work. Perhaps because a third of these AI agents are deployed without IT, 63% of IT leaders felt their team knew less "about how AI agents work" than other employees in their organization.

Trying to prevent or stop shadow AI is a futile endeavor. Instead, the focus should be on facilitating democratized AI adoption by providing business teams and individuals with best practices, support, and governance.

How AI Agents Enter Your Organization



CHALLENGE 2

I don't know what's happening in my own organization.

IT leaders are operating without clear oversight across multiple dimensions. The unknowns are widespread, from the data AI agents ingest to the workflows they support and the decisions made with their help.

Agent behavior: Forty-two percent of IT leaders lack visibility into what their AI agents are actually doing across their systems.

Agent count: One in five (21%) lack confidence in their ability to provide their CEO with an exact number of AI agents currently operating across the company.

Agent spend: Only 42% can fully account for the ROI of AI agent spend if their CFO asked them tomorrow.

Agent approval status: Sixty-two percent agree (29% strongly and 33% somewhat) that most of the AI agents currently in use at their organization were **not** formally approved by IT.

When IT leaders lose oversight at this level and across all dimensions, governance falls apart, and monitoring becomes impossible. What leaders end up with is rapidly advancing technology operating outside IT's control, even as they remain accountable for its failures.

The root cause is often that the governance came too late. IT leaders didn't prioritize governance until after AI agents were deployed and usage was spreading throughout the organization.

42%

can fully account for the ROI of AI agent spend if their CFO asked them tomorrow.

62%

agree (29% strongly and 33% somewhat) that most of the AI agents currently in use at their organization were not formally approved by IT.

21%

lack confidence in their ability to provide their CEO with an exact number of AI agents currently operating across the company.

CHALLENGE 3

I'm being held accountable for systems I can't control.

For many of the IT leaders questioned, these growing issues with current AI agent deployment and governance within their environments can be linked directly to another major concern: accountability.

Who gets held responsible?

Accountability for AI failures across the business rests with IT leaders, despite a lack of direct decision-making control or visibility into business-led deployments. Almost three-quarters (74%) of respondents said IT leadership would be held accountable if an AI agent took an action that violated policy, exposed data, or harmed a customer. That's more than twice as many as those who said accountability would lie with the AI vendor.

Nobody wants their job security tied to a technology they don't fully control or completely trust. IT leaders themselves currently have the greatest exposure, not the AI vendors supplying the agents.

Furthermore, around one in four (23%) feel they would personally be expected to resign or be removed should an agent cause a major public security or compliance incident in their organization tomorrow, highlighting just how high the personal toll could potentially be under the current status quo.

We asked IT leaders to select all the consequences they believed would follow if an AI agent caused a major security or compliance incident tomorrow. Their responses showed the breadth of fallout.

Forty-six percent said the board would demand an external audit, while 36% said the organization would attempt to handle it quietly, without disclosure.

Stop fighting distributed deployment. Instead, establish policies and hold people to them. When someone deploys an agent outside the framework, that's a documented violation. Accountability flows to them, not IT.

74%

say IT leadership would be held accountable if an AI agent took an action that violated policy, exposed data, or harmed a customer.

Nobody wants their job security tied to a technology they don't fully control or completely trust.

CHALLENGE 4

Our CEO and board are misaligned, and external pressure is making it worse.

External voices are reshaping C-suite conversations about AI. When Jensen Huang warns CEOs not to let one vendor own their AI strategy, and McKinsey analysts or OpenAI leadership emphasize the speed of AI adoption, these statements carry weight at the boardroom table.

When your CEO hears they need an “AI strategy,” they may push for faster adoption. Yet this external pressure is creating an internal fracture.

The fragmentation extends to the board itself.

When asked what would happen if an AI agent caused a major incident, responses scattered across competing strategies with no clear consensus: 46% believe the board would demand an external audit, 36% say their organization would attempt to handle it quietly without disclosure, and 24% say it would publicly blame the AI agent vendor.

Over one-third of IT leaders say their organization would attempt to handle an AI-related incident quietly, without disclosure.

This division reveals a fundamental lack of shared playbook at the highest levels, exactly when unified decision-making is needed most.

Nearly seven in ten IT leaders shared that high-profile public statements from AI leaders influence their CEO’s thinking about AI adoption.

When asked, “Which of the following best describes the conversations you are currently having, if at all, with your CEO about AI agents?”

21%

said they’re pushing for faster adoption than their CEO is ready for.

17%

have CEOs pushing them faster than they’re comfortable with.

5%

reported not having had significant discussions about AI agents.

CHALLENGE 5

I'm desperate for observability and auditability.

IT leaders expect a major AI-related public security or compliance incident. That expectation is driving a demand for traceability, accountability, auditability, and control over all agents within their environment.

Sixty-two percent of IT leaders said they would bet that an AI agent will cause a public security or compliance incident at their company within the next 18 months. Sixty percent worry more about AI agents than traditional cyberattacks. Yet they're not recommending their organizations stop deploying agents. They know that with proper governance, these tools are transformational. Instead, they're asking for better visibility into agents that are already running.


The problem is architectural. Traditional security tools (firewalls, network monitoring) defend against external threats, but AI agents running on an employee's device under their login appear as normal activity to these security tools. Existing defenses have limited visibility here.

This sentiment is strongly reflected in responses to the final survey question, which asked what priorities IT leaders most want to claw back regarding agentic AI adoption.

The top answer across all respondents was "visibility into what AI agents are actually doing across our systems" (42%, rising to 48% among CTOs), closely followed by "the ability to audit and reverse an agent action" (31%), "knowing exactly who is accountable" (29%), and "confidence that every action is compliant by default" (27%).

These priorities center on control and transparency rather than faster adoption capabilities or improved agent tools. This ranking shows where IT leaders see the real gaps.

What Control Matters Most



What this means

Our research reveals a clear gap between the rapid, sometimes rushed, adoption of agentic AI currently taking place and IT leaders' ability to govern it effectively. Unapproved agent use, limited visibility, knowledge gaps, and uncertainty over accountability are leaving many leaders feeling increasingly exposed. In fact, over a third of IT leaders are losing sleep over AI agent activity; just over two-thirds are checking systems outside working hours. Disagreements at the board level over how best to roll out the technology are adding further pressure.

The challenge now is to establish the visibility, governance, auditability, and control needed to deploy AI agents with confidence. In doing so, organizations can rebuild lost trust, reduce risk, and create a stronger foundation for responsible and sustainable agentic AI adoption both now and in the future.

Many organizations are responding by centralizing AI agent governance. In our survey, almost half of IT leaders are currently working to change their integration and orchestration strategy to prepare for autonomous AI agents, and 45% have added AI agents to board-level discussions.

Instead of distributed approval processes, they're building a single orchestration layer that provides visibility, audit trails, and compliance enforcement. This approach gives IT leadership the control they need while allowing business units and employees to deploy agents within defined guardrails.

The core issue is organizational clarity:
Who decides about AI agents? Who approves them? Who's accountable?

45%

have added AI agents to board-level discussions.

An organization needs defined answers to these questions to address the accountability gap. A centralized governance model offers a concrete path forward.

Methodology

Censuswide conducted this research among a sample of 1,001 US IT decision-makers (VP-level and above, 251 CIOs, 665 C-level executives) at organizations with 1,000+ employees.

The data was collected between May 15–22, 2026. Censuswide is a member of the Market Research Society (MRS) and the British Polling Council (BPC), and a signatory of the Global Data Quality Pledge. We adhere to the MRS Code of Conduct and ESOMAR principles.

