pendo

The CIO's guide to reducing AI-related risk

Contents

Introduction

AI is inevitable. But its potential comes with serious risks.

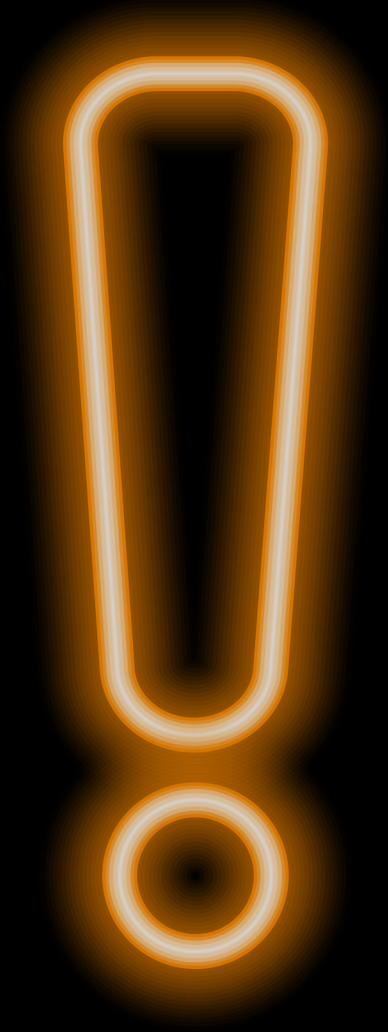

AI, large language models (LLMs), and autonomous agents are already transforming business. We're entering an era in which knowledge workers spend less time on administrative and repetitive tasks and more time creating value for the company.

Business leaders see the potential, and they're all in. A PwC survey found that nearly half of Fortune 1000 companies have fully integrated AI into their products.

The potential increases in productivity and efficiency are massive—but so are the potential compliance, security, and operational risks. CIOs must strike a balance between moving too fast and inviting disaster and moving too slow and falling behind their competitors.

Traditional governance models are too rigid, risk-averse and committee-driven to keep up with the pace of AI innovation. You need visibility, flexibility, and control through tools that evolve as fast as your tech stack does.

This guide explores the key risks that AI introduces and how to mitigate and minimize them. We'll show how companies are putting these ideas into practice right now.

Here's how to make sure your AI implementation is transformative, not destructive.

Section 1

Agentic workflows bring new challenges

The latest AI disruption is agentic workflows: the use of autonomous AI agents to accomplish tasks without human intervention. The hype for agentic workflows is at a fever pitch. In a poll of 1,500 IT leaders, 96% said they plan to increase their use of AI agents over the coming year. These investments are a sizable chunk of projected IT budgets: Another global survey found that a third of businesses expect to spend $25 million or more on AI in 2025.

There's no denying that AI agents can be incredibly useful. They can respond to customer queries, create reports, and run internal processes by themselves. This means your teams can offload repetitive tasks to AI, even complex workflows executed across multiple applications.

But the autonomy that makes agentic workflows so attractive also introduces new risks. AI agents could access and use sensitive data in noncompliant ways, such as exposing information to individuals or parties that don't have permission to access it. Agents could also give incorrect information that could put your business in ethical or even legal trouble.

For example, Air Canada's chatbot invented a coupon in the course of a chat. When the airline refused to honor the discount, the traveler sued—and won. The court ruled that Air Canada was legally responsible for what its chatbot told customers.

Section 2

Four key AI risk vectors

As AI becomes more embedded into daily operations, CIOs must be forewarned and fully prepared to address new risks.

These fall into four broad categories:

1. Security vulnerabilities
2. Compliance failures
3. "Shadow AI"
4. Software sprawl and experience erosion

CIOs are used to managing risks like these, but autonomous agents require more oversight than the average SaaS solution. Governing them will require continuous, contextual insight into their behavior.

1. Security vulnerabilities

AI agents can work both within and across applications, which makes them prime targets for exploitation. They are usually non-deterministic, which makes it difficult to apply traditional security threat modeling or design review principles. Without sufficient oversight, they could be used to export data, make untraceable changes, and even grant admin privileges to bad actors.

2. Compliance failures

Agents can be blind to compliance-required steps in workflows. They might mishandle PII, violating your customers' trust and potentially breaking privacy laws. They could fail to document processes that require an electronic paper trail, or discard data that your company is legally obligated to retain.

3. "Shadow AI"

In the same way employees use unauthorized software, teams can now spin up AI tools without IT oversight, creating fragmented and ungovernable systems.

4. Software sprawl and experience erosion

User experience can degrade as AI tools proliferate. Overlapping tools and unclear workflows frustrate employees and undermine productivity while potentially wasting precious IT budget.

Section 3

Why traditional governance falls short

Most organizations will try to manage these risks with time-tested methods that work for other types of software. You might order more employee training, limit access to tools, and draft pages of usage policies. But while these efforts are a necessary part of legal and regulatory compliance for employees, they are insufficient for agentic workflows.

While you can certainly train an AI agent with your security awareness documentation, there's no guarantee that the training will be effective. Policies can't cover every contingency. Manual audits can't keep up with behavior that can autonomously shift moment-to-moment.

Agentic workflows require oversight that is:

- **Real-time.** Know how your agents are performing and interacting with users.
- **Behavioral.** Monitor actions, not just outcomes.
- **Actionable.** Have the capability to intervene and improve.

Pendo has a proven track record of helping CIOs reduce software-related risks. Here's how we're helping empower companies to meet the challenges that the agentic era brings.

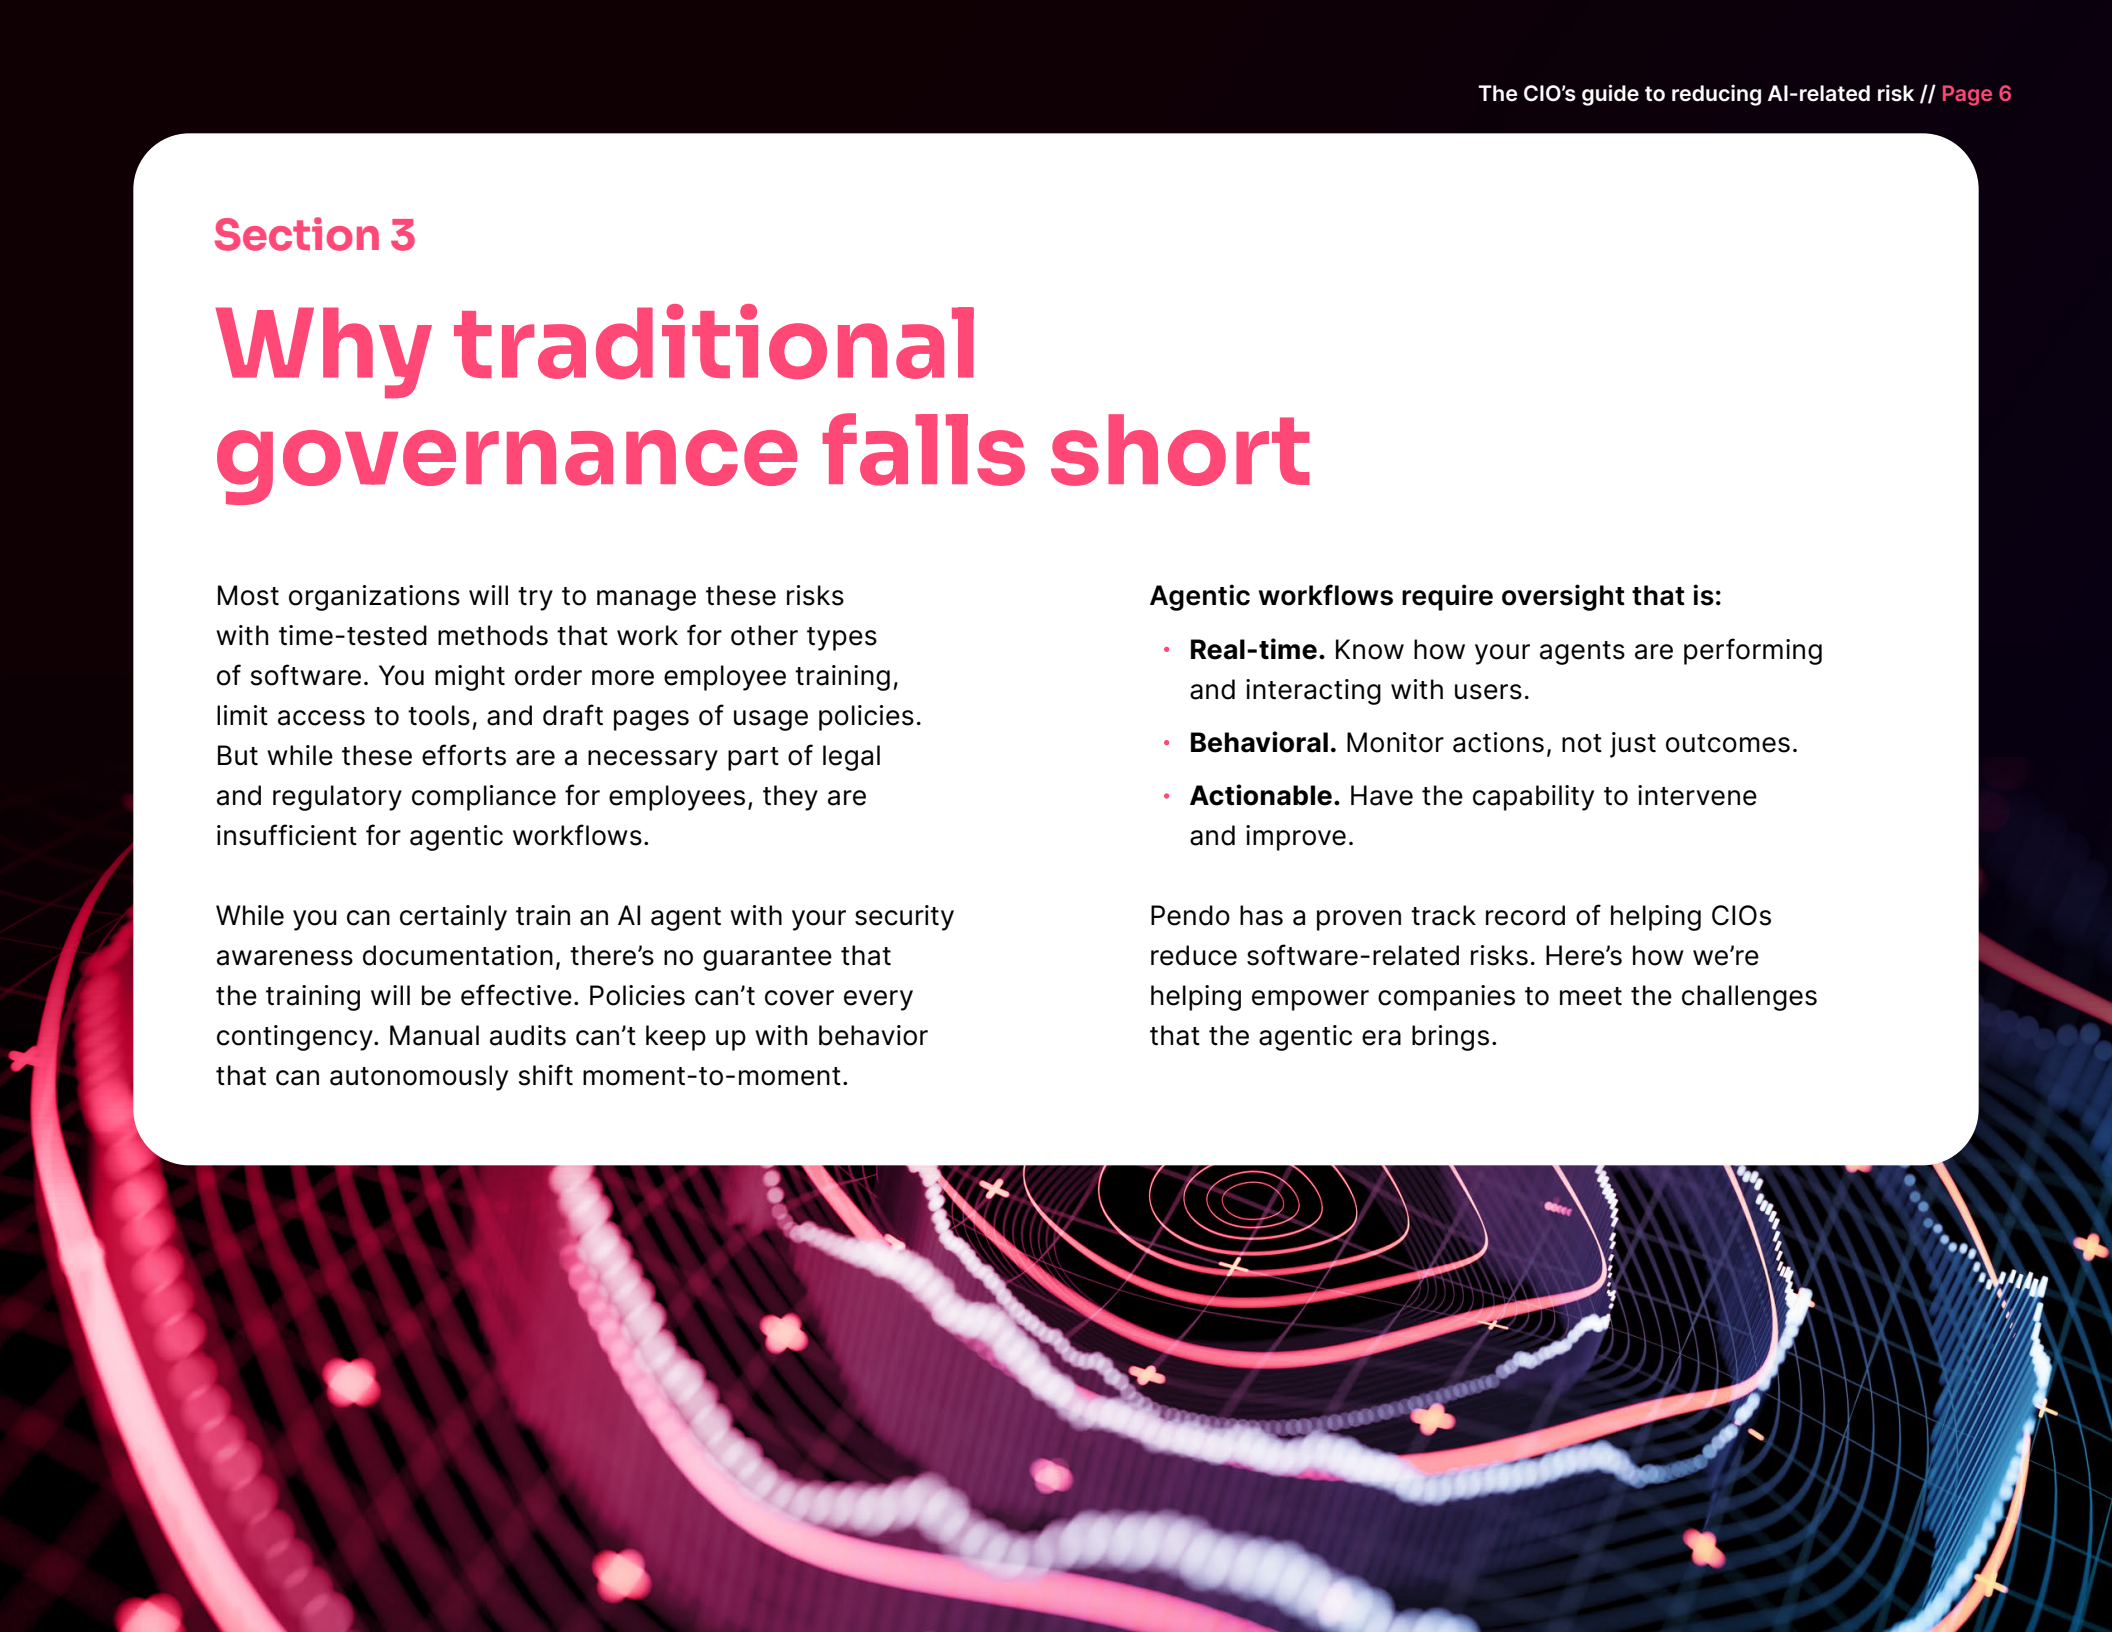

Section 4

An analytical approach to AI risk mitigation

What if you could analyze your AI agents in the same quantitative, data-informed way that you do user activity? What if you could dig into how your employees are using AI agents, identify risky behavior, and apply guides to mitigate it?

Pendo Agent Analytics is built to analyze the behavior of agents you build and agents you buy, to make sure they're driving productivity and revenue without introducing risk. It's easy to see how users are interacting with AI agents—whether users are getting value from these AI helpers, what they're being used for, and whether they're driving value.

For agents you build, Pendo Agent Analytics helps increase adoption, improve your AI roadmap decisions, and directly connect your engineering work on AI agents to business outcomes. For agents you buy, Pendo can also help drive adoption, as well as improve productivity, understand your organization's AI usage, and reduce risks.

Pendo is designed to help you unlock the full value of your software. Now, that functionality extends to agentic workflows, too.

CUSTOMER STORY

Thomson Reuters' tax and accounting division rolled out an AI-powered search engine for their Checkpoint Edge platform. They needed to understand how people were interacting with the new functionality to improve the search engine's relevancy.

They used Pendo to get quantitative insights about the search experience and then correlated the findings with NPS scores to get a complete picture. This data empowered their development team to improve their AI capabilities and ensure users got the most relevant content from each search.

Section 5

Proving ROI with Pendo

Wasted spend is a major risk for any software purchase, and AI solutions are no exception. If your organization has a fragmented tech stack with dozens of AI applications, just knowing what is and isn't getting used can be a challenge.

Pendo makes it easier to answer the burning questions CIOs have about their teams' agentic adoption:

- Are agents saving users time on key workflows and processes?
- Are agents cutting costs and helping increase revenue?
- Do they free up teams to work on higher-value tasks?
- Are they increasing outputs like task completions and conversions?

Pendo's AI-driven analytics offer true visibility into each of these questions and the tools to act on these insights, all in one seamless solution.

CUSTOMER STORY

Demandbase, a leading go-to-market platform, integrated AI agents into their platform to help customers work faster and close more pipeline. They needed to observe how users were interacting with the agents and collect in-context feedback.

Pendo made it easy to get the insights Demandbase needed.

"**We're getting input directly inside the product, right as users interact with the agents,**" said Chad Holdorf, VP of Product at Demandbase. "**We're learning whether it's helping them do their job, whether it's actually moving pipeline. That kind of feedback is critical to building the right agents and improving them fast.**"

pendo

Conclusion

Adopting AI intelligently

It's clear that AI is more than a phase or a "next big thing." It may very well be a foundational shift in how your organization operates. CIOs are tasked to lead their organizations through the AI transformation smartly and securely.

Pendo is ready to help you navigate the agentic era. With Pendo Agent Analytics and the rest of the Pendo platform, you can gain visibility into both agent behavior and how your users interact with the technology and take the actions to optimize both.

Ultimately, Pendo helps reduce AI-related risks and maximize the potential that agentic workflows can bring to your organization.

Sign up today and start using Pendo for free.