

CIOs rethink the balance between AI oversight and innovation

CIOs face mounting pressure to support AI adoption across at speed — but at what risk? IT leaders shed light on how they're balancing oversight and innovation.

BY ROSALYN PAGE · CONTRIBUTING WRITER

The new CIO mandate

The new CIO mandate is clear: facilitate AI adoption across the enterprise at speed.

According to CIO.com's State of the CIO survey, CEOs' top priority for their IT executives is to capitalize on AI. From researching to evaluating AI products, CIOs are now the central figures in their organizations' AI strategies.

And company leaders are looking for real outcomes. Almost two-thirds of senior leaders report there is more pressure to prove ROI on their AI investments than a year ago, according to Kyndryl's 2025 Readiness Report.

Numerous sources — from the board, to the CEO, to business units and competitors — are behind this pressure, says Jonathan Tushman, chief AI officer and CTO at Hi Marley, a customer conversational platform for the property and casualty insurance industry.

Succeeding in the task ahead of them requires complex conversations, and getting through legal, compliance, and other checks "at a reasonable clip," adds Tushman, who added CAIO to his remit more than 18 months ago but has felt added urgency in the past six months. In professional gatherings, board conversations, and almost everywhere across the business world, the conversation turns to AI — and then quickly the fear of failing behind.

That includes employees as well. "It's the engineering team and there's everybody else — marketing, sales, finance. It's people who are not AI-native, but they're very eager to use these tools at an early level," he says.

As CIOs find themselves facing pressure to scale and demonstrate real value, the challenge is keeping up with risk considerations — without creating unnecessary friction.

AI is a step-change in risk management

Most IT leaders are a long way from feeling comfortable with the new AI risk management balancing act. Just 31% of respondents feel completely ready across external business risks, Kyndryl's survey reports.

Tushman believes two things are genuinely different about the risks AI introduces. The first is that AI is indeterminate, whereas most technology is deterministic. "You can't prove an AI system will or won't do X, so the traditional 'put controls around it and verify' model breaks down," he says. "We need a different way to govern something whose behavior you fundamentally can't pin down."

The second is the gravitational pull on end-users. "With most tech, IT could take its time evaluating before rollout," he says. "With AI, if you don't put powerful tools in front of people fast, they'll route around you — and shadow use creates more risk than controlled access ever would. The timeline compresses at the same time the control model gets harder."

Tony Vizza, founder and managing partner of Novera, agrees that the instinct to move fast can lead to the exact failures everyone fears.

AI, risk and organisational design

“This might be staff putting sensitive information into public tools without a proper governance structure, or people copying and pasting straight out of AI and sending incorrect deliverables to customers,” says Vizza.

Organizations should avoid jumping into AI because of the fear of missing out without first clarifying where and how it will be used. All risk decisions should flow from these questions, he says. “What problems are you trying to solve — is it better customer service or deeper insight into your data? What are you actually trying to do?”

Vizza recommends guiding AI decisions with a risk assessment that considers expected outcomes, size of investment, and its importance to the organization’s objectives. “You define your risk appetite, build a risk register, and define what risk treatment should be for each risk,” he says.

“For example, if you’re going to use a public AI model, you might treat that risk by not putting sensitive data in or buying the right license so that if you do, you’re covered, or getting guidance from the regulator before you proceed.”

Organizations must also consider AI services as a third-party risk, and not leave all accountability with AI providers, Vizza says. “You can’t outsource the responsibility,” he adds.

Due diligence is required to understand what is in the AI provider’s contract, who is responsible if they have a data breach, and how your organization can pursue them if something goes wrong.

“Some organizations build that into their risk management process. Others are quite flippant or don’t even know they should be asking those questions — and that’s what gets them stuck down the track,” he says.

The importance of organizational design

At Hi Marley, Tushman and team have made structural decisions to foster “healthy internal tensions” that are intended to surface and address AI risk considerations. This includes separation between the “AI adopters” in the product and technical teams and the “AI oversight” teams in compliance and legal. Compliance owns the audits, security concerns, and ongoing oversight, while legal owns the documentation that describes the boundaries.

“The key is that it’s independent from the teams pushing AI forward,” he says.

“Companies need to invest seriously in these compliance functions. Hire smart, nuanced people. These roles can’t just be ‘no’ machines, but they can’t rubber-stamp everything either. The value is in the judgment,” he says.

Tushman’s role is the AI innovation steward, spearheading AI adoption that includes being challenged on risk, compliance, and legal considerations. “We have a senior leadership team and we have ‘conflict by design’ within that group,” he says. “I play the CAIO role and next to me, I have our head of legal and our head of compliance. So in that leadership team, if we have ‘conflict,’ we’re able to understand the trade-offs and make a decision as a group.”

Tushman believes this creates healthy tension: Innovation-minded leaders push boundaries while compliance and risk leaders counterbalance them. But if a decision can’t be reached, it goes to the CEO. “I do recommend a [split decision] goes to another officer in the organization,” he says.

Decisions about organizational structure could prove to be as consequential as the AI adoption decisions themselves, Tushman says. “The companies that get the organizational design right early will have a real advantage,” he explains.

Desire for AI advances the risk equation

One of the features of the AI wave is the thirst for access — from the board to employees — to use the tools, build applications, and start putting them to work. “Right now, everyone’s dying to try it,” says Tushman.

Hi Marley is in the “activation” phase — meeting the appetite for the tools with safety wrappers. “My main goal here is to have people learn the tools, start using them, and gain some competency with them,” he says. “We will get to the measurement phase, but I think spending too much time on measuring right now is not worth the effort.”

Tushman, like many, is watching how quickly models improve. “AI has huge implications for how you organize, how you hire, and what buy-versus-build decisions you make,” he says.

AI, risk and organisational design

Zuora, which specializes in software for subscription and recurring revenue businesses, is three years into its AI journey. Chakkarapani is adamant that speed for speed's sake is not the goal.

“We don’t want to take an existing process and just make it faster. You’re just making a process more chaotic. Can we make it fast, smarter, and reorganize it?”

Vizza believes a good percentage of CIOs will need external help to navigate the push for rapid AI adoption. “Or they’ll need to upskill themselves, because AI operates very differently to traditional IT,” he says.

His advice is threefold. First, “make your decisions on the right basis — either learn how AI really works or bring in someone who can advise you properly,” he says. Second, bring it back to the business purpose. “There are opportunities with AI, but the core question is, ‘What are we trying to achieve by bringing this in?’” And third, work out how you’re going to manage the risk.

“Risk isn’t necessarily a bad thing — Formula 1 cars are risky, but they have very good braking systems so they can go faster,” he says. “It’s the same with AI: You put the right risk management in place so the business can move quickly without suffering adverse consequences.”

In its almost three-year AI journey, Zuora started with experimentation before moving 12 enterprise-wide pilots into production, Chakkarapani says, adding that there are three pillars to assess potential AI projects against: effort, value, and confidence. “Effort includes the security risk,” he says. “Is it low, medium, or high?”

Chakkarapani’s team started with simple executions, although the first experiments didn’t go as hoped — providing valuable lessons for the following ones. “We learned AI is only good when you have the right data — the right content, context, and governance,” he says.

They moved on to IT service management and that’s when the practical learnings really started, gaining feedback from internal teams and users, answering the security and governance questions, and iterating as they went.

Early applications include marketing, sales, product, and technology, achieving 10x to 25x throughput improvements. Success is measured in business outcomes such as growth, cost saving, customer engagement.

Through this process, the team has been doing the “behind the scenes” work to speed AI adoption across the company. “We realized that to go at speed and scale, we need to have the right trust, security, and governance underlying it,” he says.

An enterprise-wide platform connects Zuora’s approved AI services, including ChatGPT and domain-specific tools, to its structured and unstructured data. On top of this is the context layer and services so that people can build their own applications. It uses each employee’s existing login and organizational profile, and it respects the same role-based security.

“We slowly developed the framework that became our blueprint with the 10 to 12 things that need to be considered when creating an AI-driven application. When someone is interested, they’re taken to the self-directed process with these do’s and don’ts that is automatically downloaded as a markdown file to that person’s computer,” he says.

The ultimate aim is delivering up to 100x business value through an enterprise-wide governed platform — covering IT, HR, finance, legal, procurement, sales, and product. IT plays the role of orchestrator, providing the platform to access the tools and agents and collaborating with the business team to reorganize that workflow.

The AI maturity model

Chakkarapani believes the more secure the environment, the more it paves the way for experimentation, adoption, and, in time, business results. At Zuora, Chakkarapani has evolved this process through three levels of organizational AI maturity to date:

Level 1: IT provides a platform and services. Employees have controlled access to data based on their role and security privileges. They can create their own agent for themselves. If something doesn’t pass the minimal security and compliance and requirements, it cannot move ahead.

AI, risk and organisational design

Level 2: An employee-built agent goes through an IT governance check for duplication or overlap, model improvements, security scans, and manual reviews. If approved, it's shared with the wider enterprise. "We're doing well on that, but it's still a lot of manual work because there are no tools in the market that can automate this," he says.

Level 3: At this stage of maturity, an organization has established a secure foundation across its applications so AI can scale safely. At Zuora, over six to eight months the team tightened endpoint and application security, enforced mobile device management, introduced AI usage monitoring (including what staff upload into prompts), and disabled Google authentication to block personal or bulk email accounts from accessing unapproved apps.

Earlier this year, the team embarked on working toward Level 4 maturity, where anyone can create a functioning application with minimal human involvement. Realistically, they expect to be 80% to 85% zero-touch because the final mile will still require human involvement.

"My goal is to provide a zero-touch service for anybody in the organization to create applications. If we do, they can go from a concept to an idea, prototype, design, and production — and they do it in less than two weeks," he says.