McKinsey & Company

McKinsey Technology

The board's role in managing emerging AI risks

Leading cybersecurity executives and board directors discuss the challenges and opportunities of AI.

May 2026

Boards are under growing pressure to oversee AI, with many actively working to identify the right tools, metrics, and expertise to do so effectively. While AI offers transformative opportunities for growth and productivity, it also introduces new and rapidly evolving risks, from security and bias to operational and reputational exposure.

During a recent panel discussion, McKinsey and the National Association of Corporate Directors (NACD) gathered top chief information security officers (CISOs) and board directors, highlighting four priorities for effective oversight: strengthening governance and accountability, balancing innovation with risk, building real-time risk-management capabilities, and improving AI fluency in the boardroom. Together, these shifts signal that AI is no longer just a technology topic; it is now a core enterprise risk and strategic differentiator (see sidebar, "On the street: Sights and sounds from the world's largest cybersecurity conference").

Moderated by McKinsey Partner Rich Isenberg, the panel included Deneen DeFiore, senior vice president and CISO at United Airlines and member of the board of directors and chair of the risk committee at Blackbaud; Guy Gecht, board chair at Logitech and director and chair of the technology committee at SolarEdge Technologies; and Manjula Talreja, independent board director serving on the boards of Verint, Imprivata, Ping Identity, and Proofpoint.

This Q&A has been edited for clarity and length.

AI changes everything, but governance fundamentals still matter

Rich Isenberg: Two years ago, most boards and CISOs did not focus on AI; they focused on cybersecurity risk, resilience, and technology change with automation and cloud. What has materially changed in the technology and risk landscape with the rise of AI, and what has not changed as much as headlines suggest? How are you and your teams approaching AI differently, if at all?

Manjula Talreja: AI innovation is moving at warp speed and needs to be delivered at token cost. Each day is a new day. But before we talk about what is happening at the board level, we must take a multigenerational view of AI's impact. Trust and humanity are the fabric of society. Children trust their parents. Employees trust their managers. Investors have trust in markets. Citizens have trust in their institutions. Now, we are being asked to trust algorithms. With trust between humans, there is something of a moral compass. AI does not have a moral compass.

As we get into "what is AI going to do?," it is embedded in every software now, whether we know it or not. Two years ago, cybersecurity and protection were all north–south, like networking. Now, AI is pulling data from everywhere, thus a major paradigm shift. This is the largest transformation of my lifetime. What has changed? Everything has changed.

Rich Isenberg: How are board members and the CISO approaching this revolution differently from other technological transformations?

‘AI innovation is moving at warp speed and needs to be delivered at token cost.’

—Manjula Talreja

Deneen DeFiore: Some people tend to say that AI is just another piece of software. It is not. There are many dimensions of risk and the uptake of technology. It is moving at an exponential rate in the ways companies are using it.

As a cybersecurity expert, I’m focused not only on AI security but also on managing digital risks across all AI areas. We have created a holistic framework across the organization, from operations, finance, legal, and risk and compliance to understanding the investments and outcomes. Is AI a value add in the current use case deployment? There is a cost versus risk balance, and AI goes beyond digital risk—it has an impact on safety in the world of air travel.

Accountability and ownership are vital with responsible AI. AI impacts operations, resiliency, revenue, and more. Therefore, CISOs must look at AI with a broader lens than just cybersecurity.

Rich Isenberg: Should boards feel sure that their current governance structures have covered recent technological risks? And where do they need to alter their oversight for this change?

Guy Gecht: The board’s role is not to manage AI. It is to ensure the company moves fast enough to capture the upside, with clear guardrails where failure is unacceptable. As board members, we need to oversee the company for the benefit of our shareholders. The speed and magnitude of changes is separating the AI winners from the AI losers. This is something that has dramatically changed.

Rich Isenberg: What are questions that no boards were thinking about two years ago on AI that they should now be asking of their executive committees?

Manjula Talreja: First, what does good governance look like? I was on the audit committee of a public board until last year. The CISO would present a heat map regarding cyber risk. With AI, those days are gone. This is a transformation, so the board and the CISO have an opportunity to gain experience and change their roles.

On the street: Sights and sounds from the world's largest cybersecurity conference

During the 2026 RSAC, more than 45,000 technology and cybersecurity experts and leaders met in San Francisco to talk about how AI is affecting companies. Three topics are top of mind for CISOs and their businesses:

1. *CISOs are not asking for more tools; they are asking for a coherent story.* Security teams want to buy systems of security, not just security capabilities**.** The current landscape does not make it simple to piece together a vision for securing an agentic world—and frankly, current vendor solutions and marketing do not help. Everyone has added agentic AI to their product, but they are still selling into silos such as identity, endpoint, data, and detection. Managing dozens of tools creates friction; layer in AI, and the complexity just multiplies.
2. *Less of 'what is happening now' and more of 'what is about to hit us.'* There is a real concern about risks. We do not fully understand them, but we know they are coming fast. Variations of the same questions come up frequently: "How do I build for risks that aren't fully defined?" "What are others seeing?" "How do I avoid being caught flat-footed?" AI adoption is driving at this fast pace. Organizations are scaling gen AI quickly, but their risk and governance models are not keeping up. This is a gap between experimentation and control. Cybersecurity is becoming much more forward-looking. It is less about protecting what you have today and more about building resilience for what could emerge tomorrow—machine-speed attacks, synthetic identities, or autonomous decision-making.
3. *AI is reshuffling who's actually mature in cybersecurity.* A challenging reality is that certain organizations previously perceived as leaders may soon find themselves lagging behind. Traditional cyber maturity was built on legacy architectures—perimeter defenses, human identities, and end points, but that does not translate well into an agentic world. At the same time, companies that are more engineering-led are pulling ahead because they're building security into how systems are designed and deployed. It's embedded in the workflow, not layered on top.

We have been through this digital disruption before with the internet, cloud, mobility, and now AI. While AI is moving much faster, organizations have learned from past digital transformations. The board should ask, "How are you changing? What's the impact on cost? What's the impact on growth? Where are humans being used?" Board questions will come automatically because this is about differentiation, sustainability, and the existence of your company. While there are plenty of questions, this is more of a board philosophy.

Balancing innovation with risk in a real-time environment

Rich Isenberg: Building on who will be AI winners and AI losers, how does the mindset of the board and the C-suite shift to be able to find the balance of defending against AI as well as using AI as a business enabler?

Deneen DeFiore: We are leaning into AI as a business enabler, and we are moving with purpose while ensuring appropriate controls are in place. But we are also managing the risks and how the board should be thinking about these risks with risk and control frameworks.

We are also investing in resilience and evolving our risk management practices to keep pace with AI adoption. Our risk frameworks work in real time and not necessarily in a "look back" analysis. This is a tremendous advantage for understanding, managing, and mitigating risks as they arise.

Guy Gecht: I think boards should look at AI across three areas: one, the use of AI internally to become more productive as a force multiplier; two, using AI in your products and services; and three, your ecosystem and where AI has shrunk that profit pool. The winners will be the companies that execute across all three.

Rich Isenberg: How are boards and management committees defining and operationalizing AI risk appetite and responsible AI, and what does acceptable AI risk look like in practice across internal productivity or customer-facing use cases?

Manjula Talreja: Each day, new aspects of risks are being defined. Companies are still experimenting, but now it is getting ready for prime time. Risk management has no substitute. With one technology company in which I am a board member, we always ask about keeping humans in the loop. We have not gone 100 percent autonomous. Even when you think you are ready for this autonomy, hallucinations happen and guardrails fail. You learn what guardrails are when you have a human factor. We are still learning the capabilities of AI, but magic and success happen when you have human–AI collaboration, oversight, and governance. We are only at the beginning of this AI journey.

Deneen DeFiore: Responsible AI requires both the agility to innovate and the discipline to manage risk at every step. This is a moment where we are going to have to change the way we analyze risk and risk appetite because it is changing every day. Companies need to establish their responsible AI framework, along with a set of principles or tenets. As mentioned previously, a philosophy should be established as to where will we use AI, where will we lean in, what will we do, and what will we not do. You also need to constantly revisit these tenets, and unlike quarterly enterprise risk management reporting, this can and will happen on a weekly or daily basis. People are more likely to accept this as a normal part of their job than an exercise to check a box before a board meeting.

Guy Gecht: This is divided among the risk surface, speed of the risk, and the responsibility or accountability. For 25 years, we protected data. Now, we need to protect judgment. Now, agents are recommending judgments to humans, which is a bigger challenge—putting guardrails in place for AI. The creation of technology committees within the board or managing committees allows a company to bring together the most talented people sitting in a smaller group and looking deeper into what the company is doing to mitigate the risk as much as possible.

Rethinking risk, controls, and resilience for AI

Rich Isenberg: At what inflection point does it shift from an innovation discussion to a core enterprise risk, and what is the different reporting that is being put in place to address this?

Deneen DeFiore: When you are talking about experimentation or lower-risk use cases, the controls and risk management practices have humans in the loop. As you get more control, you need different controls and constant and active monitoring to understand what is happening. With increased autonomy comes increased risk.

Additionally, as many organizations are still piloting AI, the industry is actively developing standards and benchmarks for what successful AI implementation looks like. What is your plan to run smoothly if there is a risk? Do you have people who are training to be able to do what you need to do? As with cybersecurity, a risk-based approach to protection is essential, but organizations must be prepared to respond. Organizations preach that cybersecurity is everyone's responsibility, but AI is, too.

Manjula Talreja: There is so much happening regarding AI risk, and organizations can lean on the vendor community for support—the very companies that develop, evaluate, pilot, and scale AI platforms.

Guy Gecht: Some companies say, "We learn from our mistakes." My recommendation is: Do not wait to learn from your own mistakes; learn from the mistakes of others. Following a well-known faulty software update by a company in 2024 that caused a global IT outage, I asked the others from the two public company boards on which I sit, "How can we avoid this from happening to us?" That created a postmortem review and analysis of our own crisis response planning. While the incident happened to another company, it resulted in a review of our own resiliency efforts.

Rich Isenberg: One thing that strikes me as an evolving challenge with AI security is identity as the new control plane as you get through this. If the board is hearing from their security teams that their organization already has strong authentication and that they know how to protect data—be very, very afraid and ask better questions.

For example, an employee may provide an agent with use permission for email or calendars. However, if you are approving invoices, more controls are required. This is analogous to a two-year-old. You may ask that child to walk down a hallway and stop at the door. However, the child may pick up a crayon, write on the walls, and knock over a table during this task. Essentially, they did what you told them to do but not exactly how you wanted them to do it.

> ‘Responsible AI requires both the agility to innovate and the discipline to manage risk at every step.’
>
> —Deneen DeFiore

So many technology operations groups are used to viewing things as good if they are up and operating. If an agent is doing the wrong thing, you need to know before the issue cascades across the organization.

Building AI-ready boards through talent, metrics, and long-term strategy

Rich Isenberg: Effective oversight means having the right talent on the board. How do boards find a balance today? Does this require upskilling, training, or bringing in new talent with AI fluency?

Manjula Talreja: It is a combination of all. It was not too long ago when boards were totally composed of chief financial officers. Today’s boards need leaders who understand governance, risk, transformation, and AI. As a board member, you should have a thirst for knowledge to understand the very topics and risks that affect your company. Some of this is learning on the job, with the right people surrounding you and speaking with AI fluency. Upskilling, training and enablement, as well as pursuing a balanced mix of talent, should all be considered by boards.

Guy Gecht: It is not realistic for each board member to be highly technical, but boards need enough collective fluency to stay current on the implications of AI.

Rich Isenberg: With so much change, how do boards and CISOs address long-term planning in the new age of AI?

Manjula Talreja: Short-term planning is centered on cost reduction, and we measure this. Last year seemed like a time of experimentation. The board asked what cost reductions we have seen from our AI implementation, and it was too early to determine. This year, we expect to have answers for our annual operating plan. This is a concrete way that the board can get involved.

> ‘Some companies say, “We learn from our mistakes.” My recommendation is: Do not wait to learn from your own mistakes; learn from the mistakes of others.’
>
> —Guy Gecht

Long-term planning is how a company will differentiate itself from its competitors. In some more mature companies, both are taking place, where investments and resources are dedicated to managing internal costs. At the same time, their boards are talking about long-term strategies such as agentic services and responsible AI.

Deneen DeFiore: There needs to be some level-setting on AI and its cost. C-suite leaders have an enthusiasm across industries for AI’s potential, and part of a CISO’s role is helping organizations set realistic expectations and prioritize effectively.

In the past, we used cloud financial operations (FinOps). Soon, we will see AI FinOps being used to manage and control both the costs for humans and AI token costs behind implementing AI at scale.

Guy Gecht: Every board needs to find a cadence where they discuss long-term planning at each quarterly board meeting. We do future casting and scenario planning of good, better, and best options, then provide management with feedback. Next, we conduct our strategy review versus cost. Finally, we track signals in the boardroom and make necessary long-term adjustments. We repeat this annually, but it is necessary to keep your long-term strategy flexible due to the evolution of emerging technologies.

Rich Isenberg: I have heard commentary on CEOs saying, “I’m seeing AI costs everywhere, except for my bottom line.” Nobody is getting money back. I hear soft metrics about productivity, but if every software developer has an extra 30 free minutes in their day and they’re going to walk their dog—well, maybe the dogs are happier, and that could be a good

thing, but it didn't return money to the bottom line, and it didn't allow the governance committees to decide if we do the same with less or if we do more with the same. These are the types of decisions that are starting to be forced.

Rich Isenberg: As we move toward full autonomy, what is the role of the board in the future?

Manjula Talreja: We have all been talking about a change in how businesses work. But this change is about how businesses work and whether or not they will survive. When e-commerce took off, this was a business disruptor where all of a sudden, the very existence of a brick-and-mortar book seller was affected. This goes back to the winners and losers, and the losers that do not adapt have two choices—disappear or thrive via an M&A approach. But this disruption is moving faster than both the internet and e-commerce phases. People, boards, and companies will evolve.

Guy Gecht: Most companies take the approach of doing more with more or doing more with the same. Companies—and their executives—get rewarded for top-line growth, not so much for cost cutting. Right now, focusing on top-line performance is the right productivity.

Deneen DeFiore: The technology is there. Now it is about integrating and operationalizing AI within your organization. The risks and barriers are trust factors, guardrails, regulatory environment, legal, operational, reputational, and so on. Building resilience in all business operations, including AI use cases, is a key part in the resilience framework. It is both a board and a management priority to understand and address areas of operational risk.

Rich Isenberg: For boards and the CISO, four things stand out. First, governance, reporting, and fundamentals still matter. Organizations must have clear accountability, a defined risk appetite, and transparent reporting and escalation. Fewer than 25 percent of boards look at cybersecurity metrics that relate to business impact. Many organizations still do not have an obvious way to measure the return on their cybersecurity investments. Second, not all AI use cases are the same. During piloting and scaling, organizations must embed appropriate risk controls, including where necessary, humans in the loop. Next, is the idea of AI fluency of the board. Boards need to understand and stay current on technology topics, asking management committees questions that will not only protect the organization, but help it succeed. Finally, AI crosses the threshold from the innovation topic into an enterprise risk category. Organizations must manage that from a reporting, metrics, approvals, and accountability standpoint.

‘AI is being measured by soft metrics, not the bottom line. If every software developer has an extra 30 free minutes in their day and they’re going to walk their dog—well, maybe the dogs are happier, but it did not return money to the bottom line.’

—Rich Isenberg

Deneen DeFiore is the senior vice president and chief information security officer at United Airlines. She is also a member of the board of directors and leads the risk committee at Blackbaud. **Guy Gecht** is board chair at Logitech and director and chair of the technology committee at SolarEdge Technologies. **Manjula Talreja** is an independent board director serving on the boards of Verint, Imprivata, Ping Identity, and Proofpoint. Charlie Lewis is a partner in McKinsey’s Connecticut office, and Rich Isenberg is a partner in the Atlanta office.

McKinsey wishes to thank Marcel Bucsescu, NACD vice president of strategic engagement, and Dylan Sandlin, NACD program manager, digital and cybersecurity content, for their collaboration and their contributions to the panel discussion.

Comments and opinions expressed by interviewees are their own and do not represent or reflect the opinions, policies, or positions of McKinsey & Company or have its endorsement.