

# AI Is Scaling Fast—How Companies Can Manage Risk

*Updated 04 May 2026*

**Steve Odland:** Welcome to C-Suite Perspectives, a signature series by The Conference Board. I'm Steve Odland from The Conference Board and the host of this podcast series. In today's conversation we're going to talk about how companies are governing AI and how are they thinking about the prioritization of risks and the rewards to their bottom line.

Joining me today is Dr. Andrew Jones, principal researcher at The Conference Board's Governance & Sustainability Center. Andrew, welcome.

**Andrew Jones:** Thank you so much, Steve. It's great to be back here once again with you.

**Steve Odland:** Dr. Jones, you've written yet another report. I know a lot of our listeners have listened to our podcasts about your reports and read them. The report is entitled *From Principles to Practice: Governing AI in the Corporation*. They'll be able to find it on our website [www.tcb.org](http://www.tcb.org) and read everything that you're going to talk about today and more.

**Andrew Jones:** That's correct, Steve. It's up on the website, would encourage all readers to look at it. This is in The Conference Board's Governance & Sustainability Center, where I sit. We're [00:01:00] really proud of this new work. It's flagship work, really. We looked at how larger companies are putting AI governance into practice.

Now, obviously this discussion, we're going to clarify what we even mean by AI governance. That included a survey of 130 senior executives, primarily at large US companies. We also had a whole series of Chatham House rule convenings last year where we dug into various aspects of AI governance and how it's becoming more

formalized and how it's evolving as AI so rapidly scales and moves across the enterprise. This is a very timely and fast-moving space.

**Steve Odland:** Yeah. It's interesting. You and I have talked about AI on a number of different fronts and we at The Conference Board have done a lot of work in all of our centers-- our Economy, Strategy & Finance Center; our Human Capital Center; our Marketing & Communications Center-- as AI really impacts and affects the various parts of the organization. You don't think about it from a governance standpoint too much. This is a different take on it. Why is governance an issue with AI right [00:02:00] now?

**Andrew Jones:** It's a critical issue because we all know AI is evolving fast. It's very quickly moving from a niche technology into the core business. We're well beyond experimental use and pilots at this point and AI is starting to touch all aspects, whether it's productivity, software development, customer-facing activities, decision-making, risk management. AI is starting to really touch strategy, controls, and reputation, and this is where governance comes in. You need good governance to manage and steer this in an effective way that doesn't amplify the risks and challenges.

**Steve Odland:** You have written extensively about a multistakeholder world-- customers, employees, owners, community, and environment are all stakeholders. Suppliers-- and how senior leadership and companies, boards of directors, need to take all of those stakeholders into account.

It is interesting that the deployment of AI also impacts all of these stakeholders. What I hear you [00:03:00] saying is this isn't just an IT issue where you implement it and it helps with productivity. You're talking about how do you govern as such that you are understanding the impact on all of these various stakeholders.

**Andrew Jones:** That's completely right, Steve. AI is so interesting in that regard because it does touch across the whole spectrum. It's becoming core to senior leadership and even the boards in a variety of ways. On the one hand, it's clearly a huge upside opportunity for productivity, for customer experience, for positioning and competition. But it's also bringing along these big potential risks, whether it's bias, hallucinations, security. It's just this dual character that really makes it so broad. To your point, it's so far beyond being a siloed tech issue at this point, if it ever was.

**Steve Odland:** If it impacts your customer, if you're doing communication, whether it's through some sort of chatbot or automated involvement with the customer, you have to have quality control and all that. Employees, the same thing, because you're dealing with people.

You're talking about assessing it from a risk [00:04:00] standpoint but this isn't let's just do an enterprise risk management. It's really the interaction and how that impacts the various stakeholders.

**Andrew Jones:** Whether you start picking up these individual parts of it or looking at it as

a whole, it raises so many big questions about who's actually responsible. To use your example, if a workforce-facing tool is implemented, is human resources (HR) responsible? Is product responsible for customer-facing chatbots? All these broader issues, of oversight, of decision-making, of escalation. It's all there in AI and it's all becoming very complex very quickly.

**Steve Odland:** Now, when AI first began to be implemented, most organizations had just a broad set of principles, make sure it's accurate. It's really evolved from a broad set of principles to a more formalized governance process. Talk about that process and how detailed does this need to be articulated in a company?

**Andrew Jones:** You mentioned principles. We found in our work that's become a kind of mainstream thing now. More than about two-thirds of our surveyed [00:05:00] companies have some kind of stated enterprise-wide AI principles. Just as recently as a year or two ago many companies were at the stage where they're coming up with these broad statements around fairness and accountability and privacy and so on. These statements matter. Many companies have them but we're starting to see the move now beyond that into we have principles-- how do we operationalize them? How do we translate those ideas into concrete decisions? Who owns AI? Who approves AI? What are the consequences when things fail? We're living through this transition where companies are starting to really move beyond that and good governance is just essential. And it's developing fast around some of these big questions of oversight and integration.

**Steve Odland:** In your conversations with these 130 or so senior executives, did they say that they were doing it granular? Are they doing handbooks? Here's the AI handbook, presumably electronically, on what to do and what not to do in a company.

**Andrew Jones:** It's a great question. As always, the devil is in the details and governance plays out differently at different [00:06:00] companies. We are seeing and we're hearing, and we saw in our survey and our convenings, best practices starting to emerge.

One I'd point to is at a lot of companies in our survey, about half, reported a centralized, cross-functional steering group. That's becoming the common anchor here. Where these groups exist, they're typically populated by IT leaders, obviously, but also legal and compliance, business units, HR, a wide sway for the C-Suite. You see in these sort of formal mechanisms companies trying to figure out how to really approach AI as an enterprise operating issue and not just a technical issue. That kind of mechanism for cross-functional collaboration is becoming key. We've seen similar things in other spaces, such as sustainability and environmental, social & governance (ESG). That's becoming a real foundation, to be sure.

**Steve Odland:** You've talked about a steering committee. There has to be some group within a company, as you're saying, whatever you call it, somebody in charge of this thing. Because AI, it's implemented by IT but it's also a business process. Implementation deals with all of these constituencies. You do need to have [00:07:00] representation across the

groups. Is there some sort of practice on who chairs this group or who's in charge and how it's run?

**Andrew Jones:** It does vary a lot. We are seeing an emerging C-Suite title of chief AI officer or an equivalent title that often is tasked with overseeing and driving this coordination mechanism. But it could also sit with the chief information officer, it could sit with the chief information security officer, it could sit with the chief technology officer. Obviously whoever has that responsibility, they're still ultimately laddering up to the board. There's a whole big piece around how the board oversees AI and how the company through management reports to it.

To answer your question, it varies. It varies by company but typically it's technology or information leaders who often have responsibility for chairing this. But not always. Sometimes it could also be the chief legal officer or even the chief ethics and privacy officer.

**Steve Odland:** Or chief strategy officer.

**Andrew Jones:** Or chief strategy Officer. Indeed.

**Steve Odland:** Yeah. It depends on the type of company, where they are in the implementation process, and where they consider the risks. Every company's different and you need to have somebody who can work across the organization and [00:08:00] is well regarded and is empowered to do this correctly.

**Andrew Jones:** That's correct. Whoever that individual or function is, it's a huge task. What are you really trying to do? We're trying to build governance around the entire AI lifecycle, right? So what is being developed or what is being purchased, what data it uses, what are the risks? Who approves it? These are the key questions that companies are grappling with in different ways and in different structures.

**Steve Odland:** What are the risks that companies are seeing rise to the top?

**Andrew Jones:** We put this to our 130 executives and we asked them to rank what are the most immediate risks to their company, and it was a long list of risks, as you can imagine. AI's touching in all kinds of ways. But a few immediate risks rose to the top. Number one by some margin was cybersecurity and data breaches. Very interesting how AI both amplifies the cyber risk and also better equips companies to handle this well.

Cybersecurity emerged as a number one, Steve, but we also see a lot of companies concerned about the legal liability and litigation exposure or [00:09:00] the evolving regulatory environment or privacy and data protection issues. There's a lot that is top of mind here for companies from a risk perspective.

**Steve Odland:** There's a reputational piece that overlays all of that. It doesn't take much with social media and the connected world to hear about some sort of mistake or some

sort of wacky thing that could possibly happen. You don't want to make headlines and you don't want to disappoint any of your key constituents. It's a list of risks but then there needs to be some sort of quality control process on top of that.

**Andrew Jones:** It's a really good point. You're so right that all of this could be looked at from a reputational perspective. You experience a cybersecurity issue-- obviously that is a material risk in terms of your systems and your data and so on. But it's also huge reputational challenge. When you look at some of the recent developments with some of the frontier general AI models, cybersecurity is top of mind for many people.

In terms of quality control and ranking, totally agree, Steve. You can't have just an endless laundry list of AI-related risks [00:10:00] in your heat map. You need to find the ways to really identify what's most material to your company, what's most pressing, and then develop the appropriate processes and controls around those, including how it all reports up to the board.

**Steve Odland:** You mentioned that the heat map, which is a key tool in an enterprise risk management process. How are companies integrating this AI risk assessment and governance process into ERM?

**Andrew Jones:** We're hearing companies are making quite quick progress in this space. It's moved fast as the recognition of AI risk has really moved with the agenda. We found about 70% of our surveyed executives said AI now appears is in the enterprise risk inventory or heat map. That's another way of saying AI has entered the formal risk architecture. That's a clear mark of maturity because AI is starting to come into the systems that the company is already using to govern risk.

The heat map's really interesting because it shouldn't just be a discussion tool. It's where companies identify risks that are significant enough to track and manage. Once AI is in that map, it can be assigned a rating, it can be compared with other risks, it can [00:11:00] be tracked, it can be discussed in regular risk review. AI, through being inserted into these systems, stops being this sort of abstract thing that might happen some point in the future and actually becomes part of everyday discussions. That's a big marker of maturity. It's the direction that many companies are taking and it's a sign of good, effective governance.

**Steve Odland:** It's an interesting thing to watch the whole digital revolution or evolution over time and all of these changes. A few years ago nobody could spell AI and it evolved to, "I can spell it now but I don't really know what it means but I don't want to fall behind." And then it was, "We're going to dump as much money as possible into this thing so we don't fall." And then it was, "We have to get a return on investment because I'm dumping too much money into this." And now it's really elevated and it's becoming part of everybody's core processes. All of that has happened in a very compressed period of time.

**Andrew Jones:** Everything you just described is what happened in about, what, three

years? Maybe four. Four at most. We've seen some big technological revolutions and we'll continue to, but in such a short space of time to go from so [00:12:00] nascent and so new to reconfiguring processes around it. And now starting to even talk about how do we reimagine companies around it. What does that mean? How do you play in this new AI-powered marketplace? It's a rapid evolution or a revolution that we're all living through the best that we can.

**Steve Odland:** Really important points. We're talking about AI governance in this modern world. We're going to take a short break and be right back.

Welcome back to C-Suite Perspectives. I'm your host Steve Odland from The Conference Board and I'm joined today by Dr. Andrew Jones, a principal researcher at The Conference Board's Governance & Sustainability Center.

Andrew, before the break, we were talking mostly about how you govern it from a management perspective and you said a couple times it that goes all the way up to the board. Let's talk a little bit more about the board of directors. What role does the board of directors play in AI governance?

**Andrew Jones:** This is a really important discussion. The board of directors is endlessly interesting. If you don't have direct exposure to the board or you don't work in the sort of the corporate governance circle, sometimes the board can [00:13:00] be somewhat of a mystery. But actually this is the group that is ultimately tasked with overseeing the company and has a fiduciary duty to do so. The way in which boards engage in AI is obviously critical to how companies proceed.

We've seen our work clear signs that board oversight of AI is evolving quick, although it's fair to say that for a lot of companies, it looks evolutionary rather than revolutionary. I don't think boards are tearing up their models and starting from scratch but they are trying to fit AI into their existing structures and adjusting those structures as they learn more about the scale and the nature of the issue.

This is a work in progress. We found from some of our benchmarking large cap company disclosures; more companies are starting to disclose where they put AI within the board structure, often putting it in the audit committee. That's starting to evolve and we're seeing more companies having a dedicated tech committee or so on.

Things are in motion right now. I don't think there's a single dominant way in which boards are getting involved but there's a clear sign of trying to get their hands around this in a more systematized and concrete way.

**Steve Odland:** It raises the question also about how much do [00:14:00] boards and the directors themselves understand AI. It's one thing if you're a sitting C-Suite person in another company dealing with these issues in real time and you've wrestled with them, you have the personal experience. A lot of boards have directors that are retired and don't have firsthand knowledge.

How are companies keeping these directors up to speed or doing director education on AI? What the risks are, deployment, all of those things. You can't just assume everybody knows all this stuff.

**Andrew Jones:** Not at all. It's a fast-moving space. We saw this play out in our survey. About half of surveyed executives perceived their board to have a moderate level of AI understanding. Only about a quarter said their board had a high level of fluency. That makes sense for the reasons you just said, Steve. Everyone is still in a process of building their AI literacy.

But what we've heard from most boards and what we've seen, most are trying to do this through a range of techniques. Some are very obvious. Having more regular board or committee discussions. Bringing in experts, whether they're [00:15:00] internal or external to the firm. We're seeing fewer companies using more structured approaches, such as tabletop exercises and gaming, or even recruiting directors for AI expertise. This is still a work in progress.

It raises big questions, doesn't it? Because what is the appropriate level of the board here? I don't think anyone wants a board that is stacked with AI engineers. They still need to be able to oversee the company and manage the company in a reasonable way.

**Steve Odland:** They're very fine people. They're very fine people, Andrew, and they're very interesting. I think you know this, right? Remember this trend, to balkanize boards, where you brought in one expert for everything. You had the cybersecurity expert, you had the merger and acquisitions experts. That model really fell apart because you need to have a board of peers and people governing as a board, not just assigning slivers to individuals.

So here we are again with another frontier. You really do need to have the entire board come up to some level of understanding and govern this together. You know that if you apply all the learning of [00:16:00] every phase before this, that's where you have to be.

**Andrew Jones:** Agreed. We've seen similar debates in the last years over all kinds of areas. Cybersecurity was a good one but also climate and I guess now geopolitics. It's not realistic to expect boards to have dedicated, deep specialism and backgrounds in all of these different areas. But directors need to obviously understand where the company is using AI, where it plans to use it, what kind of decisions AI is influencing, and what warning signs they need to keep an eye on. A fluent board needs to be able to ask management uncomfortable questions.

**Steve Odland:** It is dependent on management to do a really good job of educating your board, teaching them what you're doing and how you're doing it. And also to bring in outside experts to say, "Yeah, they're doing the right thing but you as a board should be thinking about this and this." They're using your advisors rather than assigning expertise to any one or two directors. That's been our learning over time.

**Andrew Jones:** Yeah, totally agree. It's a sort of mix of things. Again, clearly upskilling the board through internal or external experts is a key learning here, particularly as the space is moving so fast that [00:17:00] three or six months from now, there's a whole new set of issues and models we have to learn about.

**Steve Odland:** Another area this happening is regulation. It was inevitable. It was the Wild West. Now you even have AI executives out there screaming for regulation, asking for it, which always boggles my mind because if you see that it's necessary, why don't you guys get together and agree as an industry? I guess that's not always possible. But the point is AI regulation is coming, whether it's industry regulation that they self adopt or government regulation. Talk about the current landscape.

**Andrew Jones:** It's fair to say that the landscape is a patchwork. That's probably understating it somewhat. There are a lot of different things that companies are navigating at once here. We're seeing voluntary frameworks emerge that are becoming influential. We're seeing binding regulatory frameworks globally and the obvious one is the EU AI Act. That's going to have big extraterritorial reach.

And then in the US-- and this isn't unique to AI; we've seen this in other areas-- this mix of state-level rules enforcement by various agencies [00:18:00] and no single comprehensive federal AI law. So pressure coming through a lot of different ways and states approaching AI in very different ways. Colorado's a good example, it has a very comprehensive AI law. Other states focus on specific aspects, whether it's AI use in employment or synthetic AI and deepfakes and so on.

It all just adds up to a very, very complex and challenging environment. Being compliant in all these different jurisdictions at once is challenging. It may even end up being impossible.

**Steve Odland:** This is where industries need to band together through their associations and try to get whatever regulation is necessary, try to agree on it, and then take that to the federal level. You just simply can't have state and local jurisdictions doing something different when you're dealing with AI implementation on a global basis, which is really what this is. It's cross border, not cross border states but cross border around the world.

There's a brave new world out here but it's something that government relations executives need to be thinking about and business leaders need to be engaging on this to make [00:19:00] sure that it protects participants but inspires innovation.

**Andrew Jones:** Agreed. It's a challenging area for all the reasons you just said. It won't be a surprise to hear that only 9% of our surveyed executives said they feel very prepared for the emerging AI regulatory environment and most view it as very fragmented and difficult to interpret. There's a big opportunity here, to your point, for companies to also work together to perhaps try and address some of that fragmentation as well.

**Steve Odland:** You can't pick up a paper today without reading about somebody being angry about their electricity rates going up. Everything is blamed on AI these days, whether that's right or not, but there are environmental issues. You've got water use and electricity use, data, land use with the expansion, and so forth. How does all of that enter the AI governance conversation?

**Andrew Jones:** This is a really interesting area and you and I obviously spoke about this on a previous podcast about data center build out and how that has an environmental impact but also maybe an environmental opportunity as well. It's interesting observing how data center build out impacts electricity demand [00:20:00] and carbon and water.

This is still less embedded in those governance discussions than discussions about security and compliance and regulations. But it is moving up the agenda. To your point, Steve, it's becoming harder to ignore and it's becoming a major public issue. In our survey, this finding surprised me a little bit, I don't know if it surprised you. We found more than half of surveyed executives expect AI's environmental footprint to become a major business and policy concern in the next few years. It speaks to the growing noise around this issue.

**Steve Odland:** It says that our listeners should put this on their list and proactively go after this thing and make sure that you've got data center providers for every company. That you are looking at this and you know that you're thinking about it as a tier one, tier two regulatory issue, added to the environmental issues. And make sure that you're doing the right thing and that you're driving the right behavior.

**Andrew Jones:** I couldn't agree more. That plays back to the discussion we had. I don't think this means suddenly boards need a dedicated AI environment discussion every quarter. But governance leaders and corporate leaders do need to really have a focus on this [00:21:00] rapid development of AI infrastructure. Data centers can create quite material issues around reputation and operation, especially for companies that have big compute demands and significant dependency on cloud providers and, at the same time, ambitious climate commitments that now might be potentially complicated by this.

Some big questions here. We're seeing environmental impact starting to move from being just a sustainability topic toward an emerging part of the AI governance discussion.

**Steve Odland:** The other thing in the newspapers every day is people are losing their jobs due to AI. Now, there is some fact to that. There are some jobs that are being replaced by AI, certainly call centers and interactions you're doing with chatbots more. More and more people are simply trying to save costs in order to invest in AI. But this is another thing to think about and your constituents include your employees and your communities and you want to make sure that you're thinking about deployment of [00:22:00] AI as it relates to people and the impact on them.

**Andrew Jones:** I couldn't agree more, Steve. It's interesting that in our survey-- this won't surprise anyone-- 80% of our surveyed executives expect big productivity gains across

the economy from AI. But the same time, 75% said they expect AI to disrupt employment and workforce structures on a large scale. It just shows this is not a speculative side issue. It's already part of how a lot of us, including corporate leaders, are looking at AI as enterprise impacts. As you said, it's a lot more complex than simple narratives of job loss. There's a lot of different things happening here.

But clearly the people side of AI governance is important. It's not just a HR matter. It does reach into how the company affects the wider environment around it: workers, communities, stakeholder and public trust. Again, this is maybe an emerging issue. How much can individual companies control the entire labor market? Obviously they can't but there's some practical questions here around stakeholders, around people, around communities that the companies do need to be attentive to.

**Steve Odland:** They have to be deliberate in their communications. [00:23:00] Employees are adults. They need to hear it. The advice is always to be honest, open, and direct about what's happening and what's not happening so that people understand.

Okay. Final word, Andrew. Looking ahead, any other issues that were raised?

**Andrew Jones:** Looking ahead at where AI risks or area issues are emerging, we already spoke about some of the environmental and people aspects but it might be worth touching on is agentic AI because agentic AI is still emerging but it's emerging but fast. Obviously agentic AI goes a lot further than standard generative AI in how it can operate autonomously. It can take actions across all kinds of tools and systems, perhaps with limited or even no human supervision. A lot of governance leaders are very concerned about how do you effectively govern AI when it's really operating in a semi- or even fully autonomous way.

We're hearing best practices emerge here. Perhaps even looking at agentic AI as a completely separate governance category and applying different types of controls and having different ways of looking at how humans should stay in the loop. It just underscores that the whole landscape continues to evolve fast and there's big opportunities [00:24:00] here but at the same time a big need for risk management, ownership, and board-level understanding.

**Steve Odland:** Alright, we're going to leave it there. Dr. Andrew Jones, thanks for being with us today.

If you go to the [tcb.org](https://www.tcb.org) website under the Governance & Sustainability Center, and look at AI governance, you should find Andrew's latest report. Thanks for being with us.

**Andrew Jones:** Thanks so much, Steve. Always a pleasure to be here with you.

**Steve Odland:** And thanks to all of you for listening to C-Suite Perspectives.

I'm Steve Odland and this series has been brought to you by The Conference Board.

## AUTHORS

---



Steve Odland  
**President & Chief  
Executive Officer**  
The Conference Board



Andrew Jones, PhD  
**Principal Researcher,  
Governance &  
Sustainability Center**  
The Conference Board

---

The Conference Board is the Member-driven think tank that delivers *Trusted Insights for What's Ahead*<sup>®</sup>. Founded in 1916, we are a nonpartisan, not-for-profit entity holding 501(c)(3) tax-exempt status in the United States.

© 2026 The Conference Board, Inc.