



Board oversight of Agentic AI

When AI moves from assistant to actor



If 2025 marked widespread adoption of generative AI, 2026 marks the rise of systems that can act and not just assist. That shift changes the risks boards need to oversee.

What started as tools or assistants that helped draft emails or analyze data is quickly turning into systems that can initiate transactions, execute workflows, and make decisions within set boundaries across the entire company. In some cases, these systems operate continuously within clear guardrails, with management monitoring the results instead of approving every step. AI is no longer just advising management. It is starting to act on the company's behalf as digital employee within defined boundaries.

As AI moves from advisor to actor, the board's role shifts from monitoring adoption to governing authority. The question is no longer whether AI is being used, but whether its delegated authority is clearly defined, aligned with strategy, and supported by controls that can monitor and validate its actions in practice.

It means the company's control environment has to evolve with technology. When AI agents begin executing decisions, governance can no longer rely solely on reviewing outputs after the fact. There must be clear accountability, limits on what agents are allowed to do, and clear escalation paths if something goes wrong. Monitoring needs to be continuous and practical with defined thresholds and the ability to intervene promptly when results fall outside expectations. With clear boundaries and active oversight, autonomy becomes an advantage rather than a source of risk.



Agentic AI creates real opportunity. It can accelerate execution, improve decision quality and reshape competitive positioning. At the same time, it may challenge assumptions about cost structures, workforce design, and the speed at which the company can operate. Boards should ensure AI is embedded in core strategy discussions, not treated as a separate technology initiative. The board plays a key role in encouraging bold AI initiatives while making sure risks are clearly understood and appropriately controlled.

Autonomy also changes the risk model. Agentic AI can be designed with varying levels of human involvement, ranging from human in the loop to human on the loop.

As reliability of Agentic AI increases, moving from human in the loop to human on the loop shifts control from pre-approval to active supervision. In simple terms, “human on the loop” means people are no longer approving every decision before it happens. Instead, they oversee how the system operates, watch for exceptions, and step in when something falls outside established limits. AI may act within defined limits, while individuals monitor performance and intervene when needed.

This reinforces the board’s role in evaluating the strength of management’s oversight framework, including whether monitoring mechanisms are effective. Are escalation triggers clear? Have override and shutdown mechanisms been identified and tested? Is someone accountable for supervising each system? As AI moves from assisting to acting, “human on the loop” oversight becomes more important. Boards should expect that any shift from traditional review of controls is intentional, risk-based, and supported by the right mix of supervision, validation, and testing. It is not enough for this to work in policy. Management should be able to demonstrate that it is working in practice as well.



How Agentic AI changes enterprise controls

The potential impact is especially significant in financial reporting.

Agentic AI can be built directly into the business processes or ERP systems. It can continuously review transactions, check accounting decisions against company policies and standards like GAAP or IFRS, monitor how controls are working, and flag significant issues as they arise.

For example, the financial reporting process relies on periodic reconciliations and reviewing what has already happened. Agentic AI makes it possible to monitor activity continuously throughout the close, consolidation, and disclosure process. It can review transactions as they occur, compare them to company policies, flag potential misstatements, identify where controls may not be working properly, and suggest corrective steps before issues become larger problems.

Instead of reacting after the fact, finance and accounting teams can address risks as they arise.

In large and complex organizations with high transaction volumes and multiple entities, this can lead to earlier risk identification, a smoother audit process, and greater confidence in the reported results. But it also means controls need to change. If AI is involved in accounting judgments, processing transactions, or preparing disclosures, internal controls over financial reporting (ICFR) need to be updated to reflect that.

Boards, especially audit committees, should make sure control design keeps pace with Agentic AI use. Actions taken by Agentic AI systems should be properly recorded, traceable, explainable, and available for review. If the company relies on AI outputs, that reliance should be supported by documented testing and clear accountability. The supporting evidence should be strong enough to stand up to regulatory review and external audit scrutiny.

However, the risk conversation cannot stop at financial reporting. Agentic AI does not just touch the general ledger. It can influence credit decisions, hiring, supply chain, operations and regulatory disclosures. In many situations, the operational or reputational risk outside financial reporting may be greater than the impact on ICFR. Whenever AI is given real decision-making authority, the same guardrails and accountability should exist.

Agent-to-agent (A2A) is an open protocol that allows AI agents to communicate with each other in a consistent and standardized way. If agentic AI systems are interacting through a structured protocol

like A2A, those interactions can be logged, authenticated, and traced more clearly. That makes it easier to see what decisions were made, what data was used, and who or what initiated an action within financial processes. However, structured communication on its own is not a control. It needs to be supported by clear governance, defined oversight, and testing within the company's internal controls over financial reporting.

Agentic AI introduces new forms of control risk as it begins influencing reporting processes. A strong governance framework, with defined authority limits, explainability controls, and appropriate "human on the loop" supervision help manage those risks and support strengthen reporting integrity.

Beyond financial reporting, third party risk increases as vendors build Agentic AI into the platforms and services companies depend on. Organizations may be using AI without appropriate governance or guardrails around data security and reliability. Boards should ensure management understands where vendors are deploying agentic AI, what company data is being shared, who owns that data, and what would happen if the company became too reliant on those providers.

As agentic AI gains access across multiple systems, the company's cybersecurity risk also changes. When systems can initiate actions across multiple platforms within the organization, the number of

potential entry points grows. Access controls, monitoring tools, testing, and incident response plans need to reflect how these autonomous systems operate. Oversight should consider risks specific to AI environments, including manipulation of models or misuse of system permissions.

Regulatory expectations are moving quickly. As agentic AI becomes more embedded in business processes, regulators are increasingly focused on clear governance, transparency, and documented controls. The cost of not considering compliance is too great, it cannot be added later. It needs to be built into the decision process around which use cases to pursue, how authority is defined, and how systems are deployed.

Agentic AI systems that carry more risk need closer and more frequent oversight. If they materially affect financial reporting, customer outcomes, pricing, credit decisions, or regulatory disclosures, they should be tested and monitored on an ongoing basis. In some cases, independent assurance may be appropriate to strengthen confidence and defensibility.

Agentic AI may act for the company, but responsibility for its decisions and the risks it creates still sits with management and the board. Giving technology authority does not reduce accountability. It increases the need for clear oversight and strong governance.

Questions for boards to consider

Strategy and competitive positioning

- 1 How is management using AI to strengthen our strategy, not just manage risks?
- 2 Is management pushing enough to pursue meaningful AI investments while still protecting long term value? Is management focusing on AI investments in the right area, aligned to strategy?
- 3 How will management recognize if AI begins to disrupt the business, and is management prepared to adjust quickly?

Governance and accountability

- 4 Does management have a clear, company-wide view of where AI systems are being used, what they are allowed to do, and who is responsible for them?
- 5 Is the approach to governing AI consistent across the organization, or does it vary by business unit? Does the company have the right skills to provide that governance? Has management clearly defined what AI systems can make decisions, when issues must be escalated, and how the company can override or shut them down if needed?

Internal controls and financial reporting

- 6 If AI affects ICFR or other key business processes, has management updated the control structure to reflect that?
- 7 Are AI-driven actions properly recorded, traceable, explainable, and documented in a way that supports internal testing and external audit?

Cybersecurity and third-party risk

- 8 How has AI changed cybersecurity risk, and are the company's monitoring and response capabilities keeping up?
- 9 When vendors use AI in the products or services the company relies on, does management understand what data is involved, how it is restricted, and how that affects enterprise-wide risk?

Regulatory and model risk

- 10 Is management prepared to demonstrate to regulators how AI systems are governed and how these systems comply with applicable rules and regulations?
- 11 For higher risk AI uses, has management tested how well the system performs, whether it introduces bias, and whether additional independent reviews are needed?

Talent and board readiness

- 12 As the company uses AI to drive efficiently, how is management thinking about the long-term impact on human resources?
- 13 Are we investing enough in our own understanding of AI to oversee it effectively?

Boards that approach Agentic AI with clear guardrails and disciplined oversight will be positioned to capture benefits of Agentic AI while protecting long term value and their reputation. Contact us to find out how.

Contact us



Matthew Johnson
KPMG AI Audit and Assurance
Leader, KPMG US
+1 470 902 3957
mpjohnson@kpmg.com



Raymond Holt
AI Solution Leader, Technology
Assurance, KPMG US
+1 571 554 9253
raymondholt@kpmg.com



Richard Knight
Principal, Assurance
KPMG US
+1 561 426 5794
raknight@kpmg.com

Some or all of the services described herein may not be permissible for KPMG audit clients and their affiliates or related entities.

Learn about us:



[kpmg.com](https://www.kpmg.com)

© 2026 KPMG LLP, a Delaware limited liability partnership and a member firm of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved.

The KPMG name and logo are trademarks used under license by the independent member firms of the KPMG global organization.

The information contained herein is of a general nature and is not intended to address the circumstances of any particular individual or entity. Although we endeavor to provide accurate and timely information, there can be no guarantee that such information is accurate as of the date it is received or that it will continue to be accurate in the future. No one should act upon such information without appropriate professional advice after a thorough examination of the particular situation.