

Board oversight of GenAI

Business values, guardrails and governance

KPMG Board Leadership Centre

Like few topics before it, generative artificial intelligence (GenAI) has dominated discussions in many C-suites over the last year. Boards are playing a crucial role in both encouraging management to accelerate the pace of exploration of GenAI and urging management to put in place appropriate policies and guardrails for the development, deployment, and use of the technology.

As the market moves from experimentation to company-wide use and potential transformation, providing effective board oversight has never been more challenging, or more important. This paper provides directors with a foundational view of GenAI and offers insights into key areas of focus and questions to ask as the board helps management prepare for the challenges and opportunities presented by GenAI.

GenAI is moving from “market buzz” to business value

GenAI continues to make headlines and attract the attention of boards and management teams. Market focus is rapidly shifting from experimentation towards seeking tangible business value with measurable financial returns. Two recent surveys^{1,2} from KPMG in the US help illustrate where many companies are now on their GenAI journey, and where they are likely to go in the next 12 months, both from a director perspective and from the perspective of C-suite executives.

A slight majority of directors (51 percent) say that their companies are actively exploring the capabilities that GenAI offers in selective pilots and proofs of concept. Nearly 20 percent of directors say that their companies (the early adopters) have started to scale GenAI broadly across their operations, and 4 percent see the technology as already core to their business operations. This trend is likely to accelerate, as two thirds of C-suite executives say that their company plans to invest more than \$50 million in GenAI over the next 12 months. This will be spent primarily on building responsible GenAI governance programs, purchasing GenAI technology, training the workforce, and enhancing customer experience.

Interestingly, the views of directors and C-suite executives surveyed appear to diverge on the primary impact they expect GenAI to have on the company: A majority of C-suite executives expect new business models (54 percent) and/or new product or revenue streams (46 percent), while most directors (69 percent) expect increased operating efficiency. There is broad agreement, however, that continuing to build trust in GenAI, and focusing on risk management processes, data quality, and cybersecurity remain critical for business value to be realised. There is also a growing recognition that GenAI will impact many enterprise risks previously on the board’s agenda.

In short, this spells busy days ahead for boards: separating hype from reality, navigating the near- and longer-term opportunities and risks to their company, anticipating the implications for strategy, and continuing to help ensure that management has in place appropriate guardrails, governance, and compliance policies and processes around GenAI. We offer the following suggestions to help boards focus and structure their oversight efforts.

Understanding the technology from a board perspective

Core to understanding the rapid rise and potential impact of GenAI is examining how this newest member of the AI technology family differs from AI models that have been used by companies for years:

- **Instant awareness.** While it took many years for prior versions of AI to be adopted by companies at scale, consumer versions of GenAI made it in a matter of months to almost every smart phone and PC. As a result, most directors, executives, and employees have already experienced what the technology can deliver, radically reducing the time to awareness and adoption at scale. In short, GenAI “brought AI from the hands of 1,000 data scientists to a billion consumers almost overnight.”³

¹ KPMG Board Leadership Centre Survey [A Boardroom Lens on Generative AI](#), March 2024

² KPMG LLP [AI Pulse Survey](#), April 2024.

³ KPMG BLC Quarterly webcast, [A Boardroom Lens on Generative AI](#), March 21, 2024.

- **Minimal investment.** Most companies do not have to invest in building their own large language models (LLMs) to deploy GenAI. These have already been built by dozens of technology companies that have invested billions of dollars on training them to read and write, using everything available on the internet and other sources. Companies still need to invest in technology and data to connect these LLMs to the business, but the bulk of the costs have been covered. Companies can easily tap into these models for a monthly subscription fee and apply them to a dizzying range of use cases simply by writing prompts (text questions) in plain English. Traditional AI may still beat the LLMs on specific tasks but must be built and trained for every company and application, which takes time, skilled data scientists and lots of proprietary data.
- **People-centric.** Most of the business benefit from GenAI is expected to come from augmenting knowledge workers, not replacing them, e.g., by off-loading a portion of the most time-consuming, least sophisticated daily tasks. This may include activities like reading numerous documents, synthesising key takeaways, drafting initial summaries, and writing computer code. Think of GenAI as a well-read junior digital assistant with huge capacity for work that still needs supervision, rather than a supercomputer that can take over any job. This can both reduce the cost of operations and increase revenue, e.g., when applied to make the sales force more effective or when integrated into company offerings to deliver more value to customers. That said, little or no benefit will materialise if people do not adopt and reinvest their freed-up time productively, which makes change management and workforce development critical to realising value.
- **Still evolving.** GenAI continues to evolve rapidly, making board oversight particularly challenging. There are at least three technology trends for boards to watch: “Stand-alone LLMs” like ChatGPT are developing ever-more powerful versions, promising step-change in functionality; “embedded models” are being rolled out at scale by integrating GenAI “at the press of a button” into popular enterprise software like Microsoft Copilot; and large action models (LAMs or agents) are the latest buzz, promising to automate repeatable actions within corporate business processes.

As powerful as GenAI can be today, there are still many tasks for which it is not well-suited, and it is still unclear what the technology’s capabilities will be tomorrow. Boards will need to stay focused and help ensure that management considers these trends and evaluates the full range of available AI technologies in shaping their technology strategy.

Generating business value with GenAI

While we are in the early stages of GenAI, the implications for business appear significant. Boards are seeking to understand what this technology means for the company—including its operations, products, services, business model, and strategy.

In terms of **driving productivity**, using GenAI at scale is fundamentally about changing what people do every day and how they work—which requires both technology and behavioural change.

GenAI is well-suited for many of the time-consuming but not always inspirational tasks done by knowledge workers today—writing, reading, synthesising, reporting, commenting, and applying structure to data. The technology can also be a powerful tool in software coding and can yield significant efficiencies in customer interfaces and call centres.

In theory, the productivity arithmetic is simple: provide a new powerful tool and training to a knowledge worker that can free up perhaps 30 percent of the time spent today, then reinvest these expensive hours on something equally or more productive for the company. Apply the same principle to someone working in sales and the company can increase revenue in some proportion to the increase in hours spent selling. Multiply these by the number of knowledge workers who can do the same, if provided with the right tools and training, and the potential benefit can be significant—likely hard to resist for many companies:

- **To develop new products and services,** GenAI can also be deployed across the entire value chain in various ways, e.g., by making the current product development process faster, to better understand customer needs, testing new offers on the GenAI models by instructing them to role-play how a customer with certain preferences would react to new features, or by building GenAI functionality into the company’s current products to make them more powerful and easier or faster to use. This is happening today in many industries—from new medicine development to faster time to market for sneakers.

Over time, beyond helping employees in their current roles, building new GenAI functionality may also enable companies to re-engineer workflows and processes, and develop new ways of doing business.

- **Capturing the value at scale** can be done through a combination of top-down and bottom-up innovation. Initially, many companies took a top-down approach, picking individual use cases to drive proof-of-concept pilots, demonstrating value, and then rolling out the solution to a larger number of people. While some pilots have proven valuable, many have fallen short, as the use cases chosen were not well-suited for the technology, or perhaps too few roles would benefit from the solution. Learning from this, some companies are exploring “use-case factories” to apply replicable process steps to rapidly churn out new applications with higher hit rates. Other companies are using a bottom-up approach, simply providing GenAI tools (with guardrails) to many employees in different functions and encouraging the technology’s use in day-to-day work.

By monitoring what appear to be the most productive uses, companies can scale what works across the organisation. To ensure they capture the value, however, employees need to adopt the tools, share ideas, and reinvest freed-up hours in more productive areas. This is often time-consuming and will test a company’s ability to drive the behavioural change required at scale.

Some are trying to combine both approaches, starting with a breakdown of where many knowledge workers spend the most time and then targeting GenAI use-case development specifically to augment these tasks.

- **Quantifying the return on investment** to roll out GenAI at scale is rapidly becoming a priority for many companies and boards. This can be challenging, starting with the difficulty in quantifying the benefit, as the productivity gain depends on how many hours can be freed-up for knowledge workers and where those hours are reinvested. Some of the benefit can be captured just by providing safe access to a stand-alone GenAI tool, without any proprietary data, which is fast and cheap; however, a significant portion of the benefit will require the tools to access proprietary data required for more value-adding work, which will take investment and time to provide.

Managing and mitigating risks

To safely deploy the technology and unlock its business benefits, management will need to manage and mitigate a range of operational, legal, regulatory, and technical risks directly driven by GenAI adoption. According to our recent survey¹, directors are currently most concerned about the reliability of GenAI-supported data/results, cybersecurity, and data privacy:

- **Inaccurate data and results.** Our board survey showed that the top concern of directors was inaccurate data, which can result from poor data quality as well as malicious factors, such as data poisoning. Inaccurate results may arise from inputting inaccurate data, or an algorithm learning something that is wrong—or producing nonsensical or false answers (“hallucinations”). Unchecked, this can put companies and their customers at risk, e.g., denying a loan to a qualified applicant due to bias in underlying algorithms. To mitigate the risk, management can take a range of steps—from data cleansing to prompting the model to disregard certain sources of information or not to respond if uncertain. Most importantly, any application not having a “human in the loop” before taking action should be avoided and all results treated as a first draft, rather than a final report.
- **Cybersecurity risk** closely follows as a concern. Because GenAI can write code, it can also write code used to hack and create more realistic and sophisticated deep fakes and phishing scams. This can increase both the quantity and the quality of threats and substantially elevate cyber risk exposure, forcing many companies to reevaluate how they address cybersecurity.
- **Data privacy risk** is another major issue for GenAI because of the increasing regulatory focus and the maze of privacy laws and regulations. User data is often stored to improve data quality. The question arises as to whether models include data that is subject to privacy regulation, regardless of whether the data belongs to the company or a third party. If so, has permission to use the data been granted? This is a particular concern in highly regulated industries, such as healthcare.

Other risks that are likely to grow as GenAI moves from experimentation to company-wide adoption include compliance, intellectual property, and reputational risks:

- **Compliance risks** are rising with the emergence of AI-specific laws and regulations.

In March, the European Parliament passed the AI Act, the first attempt to regulate AI internationally.⁴ Now in the final stages of adoption, the European Union (EU) AI Act has broad, extraterritorial reach, and covers any entity using an AI system in the EU with significant penalties for violation. Uses deemed to pose an “unacceptable” level of risk, such as biometric categorising and behavioural manipulation, are banned. Other uses are placed within a risk tier, from high to low, with corresponding levels of compliance obligations, including additional transparency requirements for GenAI. There is no similar legislation in the US as yet; however, in October 2023, President Biden signed an executive order, “Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence,” which could be the basis for future regulation.⁵ While nonbinding, the executive order identifies eight principles to govern the use of AI, and encourages federal agencies to use their authorities to regulate AI to protect Americans from critical risks. By contrast, the UK Government has adopted a less prescriptive cross-sector and outcome-based framework for regulating AI, underpinned by five core principles. These are safety, security and robustness, appropriate transparency and explainability, fairness, accountability and governance, and contestability and redress. The UK recognises that legislative action will ultimately be necessary, particularly with regard to General Purpose AI systems, however, it maintains doing so now would be premature, and that the risks and challenges associated with AI, regulatory gaps, and the best way to address them, must be better understood. Monitoring and complying with evolving legislation and regulation should be a priority for board oversight.

- **Intellectual property (IP) risks** include the unintended disclosure of sensitive or proprietary company information to an open GenAI system by an employee, as well as unintended use of third-party IP. IP infringement is an area where there are many issues and a lack of clarity. For example, litigation is pending regarding the use of third-party data to train models and whether the use of that data for training infringes upon the third-party IP rights such as copyright. This also creates a lack of clarity about the ownership of the IP generated by GenAI models.
- **Reputational risks** are a key consideration when deciding how to develop and deploy AI at scale. Companies should develop a responsible use policy to manage risks that GenAI may pose to individuals, organisations, and society. The US National Institute of Standards and Technology (NIST) has published an AI Risk Management Framework, intended to help organisations address the design, development, deployment, use, and evaluation of AI systems to increase trustworthiness.⁶ Management teams should also consider updating the company’s code of conduct accordingly.

⁴ Text of Artificial Intelligence Act adopted March 13, 2024, by the European Parliament.

⁵ White House Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence (14110), October 30, 2023.

⁶ NIST AI Risk Management Framework, January 26, 2023.

There are also transformational risks to consider that can disrupt or delay company adoption of GenAI. Foremost for the surveyed directors are people-related risks such as talent gaps, retraining, transformation, and cultural change. As companies gear up to deploy GenAI at scale, there will be increased demand for technology professionals with AI-related skills such as model development, algorithmic development, and ensuring data quality. At the same time, companies will need nontechnology professionals who are adaptable and willing to continuously learn and upskill to use GenAI tools effectively. Companies will also need legal and compliance professionals who can navigate the ethical and legal/regulatory compliance implications of GenAI and ensure responsible and unbiased use of the technology.

GenAI-related risks identified by management should be considered on a risk scale that enables management to evaluate whether those that pose the highest risks are worth taking. The board should understand who in management identifies, evaluates, manages, mitigates, and monitors these risks; whether there is a management point-person for GenAI; and how frequently assessments are updated. The board should allocate sufficient time on the board agenda with the right members of management to ensure open communications and effective oversight. Has management engaged in scenario planning to understand the magnitude and potential interdependencies of these risks? Does management have the right governance structure and leadership in place to manage the range of risks posed by GenAI?

Guardrails and governance: Practical considerations for board oversight

With GenAI affecting multiple aspects of a business—strategy, risk, ethics and compliance, talent, human resources, operations, brand, and reputation—a broad range of C-suite functions may be involved or have responsibility and accountability for various aspects of GenAI (see Who is on point for GenAI?). This highlights the challenges and complexity of GenAI adoption and use, as well as the need for leadership and coordination at the most senior levels of management.

Given the potential strategic importance of GenAI and the complexities and risks associated with the technology, it is critical that the board focus on management's policies for the development, deployment, and use of GenAI. Key topics to be addressed in management's policies include:

- Determining how and when a GenAI system or model—including a third-party model—is to be developed and deployed, and who makes that decision
- Maintaining an inventory of where GenAI is used
- Designating a management point-person and a cross-functional team with responsibilities for GenAI
- Establishing responsible GenAI use policies that align with the company's values and address ethical issues and legal compliance
- Managing, mitigating, monitoring, and reporting on the risks posed by GenAI

- Staying apprised of the rapidly evolving regulatory landscape and ensuring compliance with applicable laws
- Monitoring and ensuring the quality of the GenAI data (inputs and results).

More broadly, the deployment of GenAI should prompt companies to take a hard look at the quality of the company's data, data governance practices, and technology capabilities. Achieving the hoped-for productivity and efficiency improvements with GenAI will depend on the quality of the company's data and how it is processed and stored. The quality and accuracy of the company's data, and how it differs from competitors' data, will be critical to competitive advantage. Boards need to have insight into whether companies are making the right investments in IT infrastructure and data quality to help ensure that the company's GenAI output is accurate.

Who is on point for GenAI?

In our recent surveys, we see leadership of GenAI often distributed between the chief executive officer (CEO) and multiple C-suite members:¹

- From our board survey, it is clear that multiple C-suite executives play significant roles in leading GenAI in addition to the CEO, e.g., the chief technology officer/chief information officer (CTO/CIO) has significant responsibility for various aspects of GenAI in three quarters of organisations, and is the most influential, aside from the CEO, in half.
- The general counsel, chief financial officer (CFO), chief risk officer, and chief operating officer (COO), also have significant responsibilities in 30 percent to 40 percent of the companies, but each is the most influential, aside from the CEO, in less than 10 percent.
- Among surveyed executives who currently point to either the CEO, CIO, or chief innovation officer as the leader for AI, the CEO is the most common lead in financial services, industrial manufacturing, healthcare, and life sciences. The CIO is on point most frequently in consumer and retail; energy and chemicals; and technology, media, and telecom.
- A growing number of organisations (about 45 percent) already have, or are considering creating, a leadership role for GenAI, e.g., chief AI officer (CAIO).

Focusing and structuring board oversight

Boards are also considering how best to oversee GenAI. Since these are early days—with the technology developing rapidly and its potential impact on the strategy and business model uncertain—oversight is largely still at the full-board level, where major issues (strategic and/or transformational) typically should be addressed.

However, some board committees, such as the audit committee or a technology or risk committee, may already be involved in overseeing specific GenAI issues.

Oversight structures will likely evolve as GenAI programs evolve. Ultimately, oversight of GenAI, like oversight of sustainability, may eventually touch all or most board committees.

Starting with an inventory of where GenAI is currently being used, boardroom conversations should focus on the reason(s) GenAI is being used, who has algorithmic accountability, whose data the algorithms are being trained on, how the company is monitoring for data bias, and how third- and fourth-party risks are being managed. Other key areas of board focus should include:

- Near- and longer-term benefits and risks to the company and its strategy posed by GenAI.
- Management's efforts to design and maintain a governance framework and policies for the development and use of GenAI—including policies to embed the guardrails, culture, and compliance practices to help drive trust and transparency in tandem with the operational or transformational benefits.
- Understanding the relative magnitude of new GenAI risks. How must prior data governance policies and processes—including data quality, IT/cloud infrastructure, and cybersecurity—be modified in light of GenAI? What GenAI risk management frameworks will management use?
- Where and how GenAI is being used by employees, and the most urgent talent gaps to fill to be competitive in the near future. Have scenarios been developed for how the workforce may need to change over time as a result of GenAI?
- Planned investments and expected returns from GenAI deployment this year, expected impact on budgets next year, and scenarios for how the financial plan will change over the next three to five years.
- Management's monitoring and compliance with the patchwork of rapidly evolving GenAI legislation/regulations.
- Securing GenAI pipelines against adversarial threats.

Issues that may require the attention of the audit committee include:

- GenAI use in the preparation and audit of financial statements and drafts of Securities and Exchange Commission and other regulatory filings
- GenAI use by internal audit and the finance organisation, and whether they have the necessary talent and skill sets
- Development and maintenance of internal controls and disclosure controls and procedures related to GenAI
- Oversight of compliance with GenAI laws and regulations
- Cybersecurity and data privacy risks associated with the use of GenAI (as many audit committees already oversee these risks).

Education and expertise in the boardroom

Another important question for boards is whether they have the knowledge, access to experts, and ongoing education to effectively oversee the company's use of GenAI. While all board members need to educate themselves about GenAI, generally, we see boards pushing back against the concept of specialist directors.

As a practical matter, it is likely that the number of individuals who are both steeped in GenAI and have the broader skill sets to be good directors is fairly limited. Boards need to consider how central GenAI is to a company when they are considering the level of director expertise required. Even if the board decides GenAI expertise is required and recruits such an expert, other directors should avoid deferring to one director as a specialist in an area that all directors need to understand and be conversant in.

Ways for directors to gain additional knowledge about GenAI include hiring outside experts, management presentations, presentations by third parties, independent reading, and talking to peers. Boards may also consider forming an advisory board to help the board get up-to-date, high-quality information on the subject from third parties.

Gauging progress of the company's GenAI journey

As significant business model implications and competitive fallout become clearer, and as broader GenAI adoption trends unfold, boards can gain a sense from management of where the company is on its GenAI journey, who is in charge, what plans are in place, and management's sense of urgency in moving forward with GenAI. The following set of questions to specific members of management may be helpful to consider:

- **To the CEO:** What are the company's aspirations for GenAI and strategy to get there? What near- and longer-term benefits and risks to the company and its strategy are posed by GenAI? Who in management is on point for driving and coordinating the GenAI transformation, and how is the work being distributed and orchestrated across multiple C-suite executives? Has management considered appointing a CAIO to spearhead the change? How do you envision incorporating GenAI into our corporate strategy process and operating goals going forward?
- **To the CFO or chief strategy officer:** Assuming that our customers, competitors, and suppliers are also rolling out GenAI, what would that do to our company's revenue and cost over the next one, three, and five years? What revenue is at risk? What new revenue can be generated? What costs will be reduced? What price pressure or opportunity does the company see?

How much has the company invested in GenAI this fiscal year, and how much will be budgeted for next year? How will GenAI be used in within your function, e.g., for the preparation and audit of financial statements and drafts of regulatory filings?
- **To the COO, CTO, and CIO:** Where is GenAI currently being used, how many employees can safely access GenAI tools at work, and how many are actively using the tools to be more productive?

Have we connected these tools to our own proprietary data? If not, will we—and when? What data are algorithms being trained on, who owns the data, and how is the company monitoring for quality and bias? What measurable productivity improvements should this translate into this year?

- **To the CSO, chief revenue officer, CMO:** How are we using GenAI to sell and deliver our current products and services more efficiently and effectively? Are we embedding GenAI into our products and services to make them more attractive to customers? What new offerings are we planning to take to market? Do we need to change our price levels or structure to capitalise on these changes? What current revenue streams are most at risk if competitors roll-out GenAI?
- **To the chief risk officer, head of enterprise risk management, or chief information security officer:** What do you see as the major AI-related risks that we need to tackle first? What GenAI governance framework and policies have we implemented already and what comes next? Are the company's guardrails and compliance practices sufficient to help drive trust and transparency in tandem with the benefits? How have we increased our cybersecurity over the last 12 months since GenAI arrived?

Answering these core questions won't be easy for most executives in the near term, but it should provide the board with a good picture of the here and now, and what the near- and longer-term future could look like with GenAI.

Balancing opportunity and risk with GenAI will be a difficult challenge. This will require directors to productively challenge management to go faster—to avoid missing out on new opportunities—while still going slowly enough to manage the risks posed by the deployment and use of this revolutionary technology.

Contributors

John H. Rodi

Leader, KPMG US Board Leadership Center

Claudia Allen

Senior Advisor, KPMG US Board Leadership Center

Per Edin

AI Lead for Advisory, KPMG LLP

Patrick Lee

Senior Advisor, KPMG US Board Leadership Center

Jeanne Johnson

Principal - Advisory, KPMG LLP

The KPMG Board Leadership Centre

The KPMG Board Leadership Centre offers support and guidance to non-executive directors, whether managing a portfolio non-executive career or embarking on a first appointment. Membership offers you a place within a community of board-level peers with access to topical and relevant seminars, invaluable resources and thought leadership, as well as lively and engaging networking opportunities. We equip you with the tools you need to be highly effective in your role, enabling you to focus on the issues that really matter to you and your business.

Learn more at www.kpmg.com/uk/blc.

Contact us

Timothy Copnell

Board Leadership Centre

T: +44 (0)7801 520802

E: tim.copnell@kpmg.co.uk

www.kpmg.com/uk/blc



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 on such information without appropriate professional advice after a thorough examination of the particular situation.

© 2024 KPMG LLP, a UK limited liability partnership and a member firm of the KPMG global organisation 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 organisation.