



The Business School
for the World®

AI Governance Principles for Boards



Contents

Foreword	03	Active technology and security oversight	11
Acknowledgements	05	Workforce transformation and human accountability	13
Introduction	06	Building trustworthy AI	15
Methodology	07	The work of the board	18
Strategic oversight for long-term value creation	08	Conclusion	21

Foreword

In my many conversations with board members across different geographies, sectors, and ownership models, I see a growing recognition that AI is now shaping how decisions are made and how value is created. Equally though, I see shared uncertainty about what it means, in practice, for board directors.

This technology brings boards closer to decision terrains they have historically approached with caution or distance. Directors are faced with overseeing strategies for technology architectures, workforce transformation, data usage, and ethical guardrails, knowing that decisions companies make now will shape business models, risk exposure, organizational culture, and long-term resilience. Those decisions will determine whether an organization thrives in the new AI world or becomes irrelevant.

In this environment, the board's role is more critical than ever, even while they are not tasked with managing AI directly. The mandate is strategic oversight: setting expectations, challenging assumptions, and encouraging ambition with a clear understanding of the shifts and trade-offs involved.

AI does not give boards the luxury of time to observe first movers, understand second-order effects, and incrementally adjust approaches. It sets instead a pace that challenges the traditional rhythms and scope

of board oversight. Boards now find that they need ways to create decision-making and oversight cycles that are fit for an environment of ongoing structural change rather than one of temporary disruption.

In examining this challenge, we were lucky to have advice and direction from many board members, investors, experienced professionals, and governance specialists. Those conversations, and a shared recognition that the conditions for effective board governance have changed fundamentally, led to the decision to follow a principles-based approach. I am grateful to these experts whose contributions helped ground the principles in real and often difficult governance conversations.

By using these AI Governance Principles as a guideline, to identify key boardroom dilemmas, structure debate and to sharpen judgment, I am confident that boards can skillfully navigate one of the most significant governance challenges of our time.



Annet Aris

Senior Affiliate Professor of Strategy,
Director of the INSEAD Corporate
Governance Centre

Foreword

We’re entering the most consequential phase of AI yet, and the boardroom is positioned to be a catalyst for real transformation.

What began as experimentation with predictive and generative AI tools has quickly evolved with the rise of AI agents that plan, act, learn and collaborate with growing autonomy. They are reshaping how decisions are made, how work is done, and how value is created. In parallel, the technology continues to advance at an unprecedented pace, and organizations must be adaptable and ready for the next wave of disruption. The workplace and workforce of tomorrow may look entirely different from today.

This shift brings powerful opportunities for growth. AI enables organizations to accelerate innovation, reimagine business models, elevate human judgment, and operate with speed, insight and resilience. For many, this will be a source of growth and competitive advantage. For some, it may feel existential. The question for boards now is how confidently and responsibly their organizations will harness AI, and how to oversee the right technology and talent strategy to enable it.

Robust corporate governance is what enables organizations to move both confidently and quickly. Trust is not a constraint on AI, but the foundation that allows AI to be scaled successfully. Organizations that

invest in governance, accountability and transparent, trustworthy AI practices will be best positioned to innovate boldly and earn lasting confidence. An evolving backdrop of AI regulation makes this even more critical.

AI’s business value ultimately depends on people. When employees trust both the technology and that their organization is deploying it responsibly, they are more likely to embrace it. Board members must also confront their own level of AI readiness. It’s imperative that boards understand the impact and potential of AI, and they should expect management does the same. The same is true of consumers and customers. The stakes are high and the decisions companies make today will likely shape the organizations that lead tomorrow. Boards can help companies with their ambitions in this time of opportunity, encouraging them to evaluate the longer-term opportunity, and to take appropriate risks to work towards it.

These AI Governance Principles respond directly to that leadership imperative and give boards a practical path forward. KPMG is grateful for the contributions of experienced board members and external advisors, and for our strategic partnership with the INSEAD

Corporate Governance Centre, whose insight and global engagement with directors has brought an incredible academic rigor and broad perspectives to this work. This collaboration underscores our shared belief that thoughtful governance can unlock, not inhibit, AI’s full potential. Boards that approach this moment with responsibility, courage, and optimism will oversee this transformation with purpose and confidence.



Steve Chase
Global Head of AI and Digital
Innovation, KPMG International and
Vice Chair — Artificial Intelligence
and Digital Innovation
KPMG in the US

General disclaimer:

The information contained in this document is of a general nature and not intended to address the circumstances of any particular individual or entity. It is provided for information purposes only and does not constitute, nor should it be regarded in any manner whatsoever, as professional advice. Although we endeavour 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.

Legal disclaimer:

This document has been prepared for general information purposes only. It is not intended to be, and does not constitute, legal advice. Readers should seek their own legal and other appropriate professional advice regarding the application of matters discussed to their particular circumstances.

Acknowledgements

The AI Governance Principles for Boards are the result of an international collaborative effort. From the outset, this project was shaped through dialogue, challenge and reflection, and a common understanding that effective AI governance cannot be developed in isolation, nor from a single perspective.

● We are especially indebted to the project Steering Committee, a group of experienced board members who collectively serve on over twenty-five company boards, and who joined the project at its inception. Their early engagement helped shape the direction of the work, surfaced real-world board dilemmas, and helped us to test the Principles against the realities of boardroom decision-making. We thank:

- | | |
|------------------------------|-------------------------------|
| • William (Bill) Coen | • Anna Gatti |
| • Monica de Virgiliis | • Christina Law |
| • Philippe Vimard | • Hélène Auriol Potier |
| • Denise Koopmans | |

● We are also grateful to our Advisory Council — investors, board members and AI specialists who serve on the boards or work directly for a further twenty-two companies and organizations. They provided valuable additional perspectives through review, challenge and feedback as the Principles evolved, strengthening the global relevance of the work and ensured that it reflects the diversity of board contexts in which they will be applied. We thank:

- | | |
|----------------------------------|-----------------------------|
| • Ann-Kristine Achleitner | • Bryan Foss |
| • Maya Bundt | • Diana Glassman |
| • Christine Chow | • Samantha Kappagoda |
| • Isabelle de Wismes | • Rory Keane |
| • Rachel Empey | • Jackie Shoback |

Our colleagues at INSEAD provided academic leadership and technical and global governance expertise. We would like to thank **Annet Aris**, Senior Affiliate Professor of Strategy, **Theodoros Evgeniou**, Professor of Technology and Business and **Tanmay Dangi**, Director, INSEAD Corporate Governance Centre, whose insights, challenge and collaboration drove the development of this work.

We also wish to acknowledge the many contributors across KPMG who shaped, tested and refined the Principles through their experience of working with boards and senior leadership teams across sectors and jurisdictions. Our thanks go to **Steve Chase, Samantha Gloede, Per Edin, Matt Bishop, Neil Morris, Rene Vader, Dhawal Jaggi, Dan Konigsburg** and **Una O’Sullivan**. Each contributed perspectives drawn from strategy, technology, risk, public affairs and board engagement, strengthening the substance and the practical applicability of the Principles.

Introduction

As a lever for transformation and growth, AI now holds the potential to affect every part of a business: from long-term strategy to tactical leadership decisions; from executives to newly hired employees; and across all functional domains.

- **Oversight in a fast-moving environment**

AI’s rapid evolution and its broad impact across the company and its surrounding ecosystem puts it at the top of the agenda for boards across the globe. Boards are tasked with overseeing how their companies transform themselves with AI while retaining hard-earned trust. They must consider how AI will impact their operating environment and the company’s ecosystem, and how the work of the Board itself must change to adapt to this new technology. All of this is happening in a fast-moving landscape where board members can only partly rely on their previous experience.

As AI moves from experimentation and pilot initiatives to enterprise-wide deployment and adoption, companies are assessing how they can create and sustain long-term value with this technology. Boards have a duty to oversee the approach by their companies to capture value from these opportunities responsibly and to mitigate against the risk of the company moving too slowly, or being disrupted, while competitors seize the advantage and take market share.

- **Balancing risks and opportunities**

Simultaneously, AI introduces new inherent risks such as hallucinations, biases, data provenance and quality issues. It can also amplify existing risks to new levels of concern, including cyber breaches and unsafe use of unauthorized AI tools by employees.

It is the board’s role to help address the dilemmas that AI adoption brings, balancing competing priorities, opportunities and risks within an environment of high uncertainty. They are challenged to govern at two speeds: overseeing short-term activities and risk controls as well as examining long-term, system-shaping choices, while retaining trust. The stakes are high because the decisions that companies make now will shape how their organizations will be transformed for years to come.

- **A principles-led, board-focused starting point**

Boards should not manage AI or replace executive judgment. They must, however, engage deliberately with the strategic and transformational dilemmas AI creates, while remaining firmly within their oversight role.

This is why our approach is principles-led. Principles are foundational beliefs that guide decision-making and behavior, unlike rules which give more precise instructions. Like the principles behind national corporate governance codes (e.g. UK, Japan, the Netherlands), they are particularly useful when confronted with situations that are unclear and uncertain.

This paper offers five AI governance principles which can serve as guidelines and as a pragmatic, logical starting point in this fast-moving world. The principles came from an examination of the dilemmas that boards face when making decisions around their companies’ AI strategy and its implementation, with inputs from over 25 board members from the Americas, Europe and Asia.

The principles focus on growth, transformation and new opportunities, as well as restraints and risk. They detail how a board, regardless of company size, sector or location, can build its AI capability, exercise its oversight responsibly, and develop processes for overseeing the present and future impact of AI on their organizations.

The principles outlined in this document are intended to support informed boardroom discussion and judgment in an area that continues to evolve rapidly. They are illustrative in nature and are not intended to establish mandatory standards of conduct or prescribe minimum governance requirements, or legally enforceable obligations for boards or directors.

Boards are invited to apply these principles in a manner that is proportionate to their organization’s size, sector, strategy, risk profile, technological maturity, and applicable legal and regulatory environment, while remaining within their respective fiduciary duties and oversight role.

Methodology

We began by convening a steering committee and an advisory council of senior board members, investors and governance experts, and AI specialists from INSEAD and KPMG. Building on evolving academic insights regarding AI governance, we brought our external advisors together for group discussions to examine real-world board dilemmas, trade-offs and complexities arising from the fast-developing nature of AI and from AI’s growing impact on board responsibilities.

As the discussions progressed, it became clear that the dilemmas and trade-offs fell naturally into three perspectives: Company, Ecosystem and practices internal to the board. Within each perspective, we discussed the issues with our advisors, exploring how these issues are considered, or could be considered, within their own board meetings.

Company perspective

What key choices should the company make to optimize long-term value creation from AI?

There is considerable uncertainty regarding the precise future use of AI. Companies must continuously make key choices about the role of AI in their strategy and operations, and about how to protect or evolve their business models and competitive advantage. The questions that AI raises around technology and workforce transformation have such a profound impact on future value creation for the company that we have included them explicitly as separate categories within the company perspective. Key dilemmas for boards to oversee are:

- **Strategic oversight:** protecting the company’s ability to create long-term value and to make investments in (AI) capability building while operating in uncharted territory where speed, experimentation and quick results are expected.

- **Active technology and security oversight:** balancing technology sovereignty, cyber-, data- and AI-security with the increased agility, speed and benefits of scale offered by partnering or outsourcing
- **Workforce transformation and human accountability:** balancing productivity gains with an effective, forward-looking workforce and talent management strategy, while preserving human judgment.

Ecosystem perspective

How does the company’s use of AI affect its various stakeholders (customers, suppliers, partners and society at large), how is AI creating external risk factors to its ecosystem, and how do these considerations operate within the fiduciary duty board members have to act in the best interests of shareholders generally? What are the

“second order” effects of a company’s AI strategy on society and how might these, negatively or positively, impact the company in the long run?

Internal board perspective

What is the impact of AI on the functioning of the Board itself, including the capabilities it now needs, its board structure and its internal processes? How can the board enable effective AI oversight of risk management at both the board and management levels, while also enabling innovation and agility?

The one-to-one interviews and group discussions with these experienced professionals helped us to prioritize the most common, complex, and impactful dilemmas. From these insights, we then crafted governance principles to facilitate board oversight for each.



Strategic oversight for long-term value creation

Principle 1

The company board is responsible for overseeing the company's development and execution of its AI strategy and how it supports sustainable long-term value creation.

- 1.1 Sustainable long-term value creation with AI
- 1.2 Responsible innovation
- 1.3 Foundational investments
- 1.4 Organizational flexibility and learning
- 1.5 Transformation scope

Strategic oversight and a focus on long-term value creation in line with the company’s mission are at the core of a board’s responsibilities and are growing in complexity with the advent of AI. In the age of AI, where many outcomes are still unknown and carry significant disruption risk, traditional approaches to strategy might not suffice.

Long-term value from AI will not derive solely from productivity gains. Boards will expect management to take a broad view of AI opportunities, e.g., to enhance decision making, growth, innovation, customer satisfaction or employee experience. Also, AI may, in many cases, change market dynamics and reshape the ecosystem of customers, suppliers and competitors.

Boards need to consider how their companies take these dynamics into account and how they might affect the company’s broader enterprise strategy and future role in the ecosystem, including consequences for the company’s business model, partnership strategy and M&A strategy. For example, an important discussion in the board room might center around the risk and benefits of more radical approaches to capture the benefits of AI.

Given the high speed of AI development, companies may choose to experiment and learn at a fast pace. In this race for competitive advantage, they risk paying too little attention to scalability, focus insufficiently on the highest value opportunities or neglect systematic learning. They also risk ignoring the unintended consequences of the introduction of AI (for example, how new capability can enable behavior or outcomes that go against their own values and the values of their stakeholders). Boards may look to see that AI development is set up in a balanced way and that sufficient time is invested in establishing appropriate guardrails.

While the benefits of AI can be significant, realizing these benefits will, in many cases, require sizeable investments in technology and data infrastructure, foundational capability building, and change management early on when returns are still uncertain or hard to quantify. At the same time, there is often pressure to show short term productivity improvements and growth.

In their oversight of resource allocation, boards face the dilemma of how to allocate AI investments between shorter-term results and longer-term innovation and transformation, ensuring a consistent and proportionate evaluation of AI costs, benefits and risks, in line with the company’s values and risk appetite.

1.1

Sustainable long-term value creation with AI

The board oversees (and in some jurisdictions sets strategic goals for) the company’s ability to create long-term value responsibly with the help of AI, balancing innovation, resilience, and stakeholder interests in a rapidly changing environment.

The board expects management to adapt its strategic approach and frameworks to reflect the unique challenges posed by AI. These include AI’s disruption potential, its high speed of development, the importance of new types of resources, and the increased need to partner and to collaborate across functions to create a fast-learning organization.

This oversight incorporates the broader ecosystem, such as vendors, distributors, competitors (including new market entrants), customers, and external advisors, to promote transparency, balanced innovation and resilience.

1.2

Responsible innovation

The board promotes a culture of responsible experimentation and innovation with AI. It monitors management as it explores how to leverage the technological advancements in AI and other transformative technologies while being guided by clearly defined values, robust risk management, clear executive accountability, and the company’s long-term interest.

The board expects management to understand the impact of innovation on how AI-enabled improvements may affect the company’s overall value proposition to customers, and how it may influence expectations around service quality, pricing, and the company’s brand and reputation.

1.3

Foundational investments

The board monitors whether the company invests sufficiently and promptly in the foundational elements of AI that enable future innovation and resilience, e.g. data, infrastructure, talent, change management, partnerships and M&A. It considers how these investments are aligned with the company's overall business model and long-term strategy.

1.4

Organizational flexibility and learning

The board considers how well the company is prepared to adapt to the rapid pace of AI development with the agility to scale up and down and potentially change course, depending on the context. It expects management to promote a culture of responsiveness, responsibility and innovation suited to the company's strategy, risk-appetite and operating context. Barriers to capturing AI potential are also considered, e.g. the need to re-engineer business processes, and ways to collaborate across functions and between humans and AI.

1.5

Transformation scope

The board expects management to consider the extent of transformation that is desirable and possible. It discusses with management the impact AI will have on the company's fundamental business and IT architecture, as well as the impact on the overall organization and the measures for success. It asks management to consider the "art of the possible", ranging from maximizing efficiencies and streamlining workflows to rethinking, where appropriate, the extent to which AI will impact the overall ecosystem the company operates in, including customer demand, business model change, and value chain composition.



AI is reshaping how decisions are made and how risk emerges. In this environment, governance must move from periodic oversight to continuous discipline. For organizations and business leaders in the Kingdom of Bahrain, the differentiator will be their ability to embed accountability, transparency, and clear decision rights around AI. Governance is no longer about control - it is what enables leadership teams to scale innovation with confidence and sustain trust with regulators and stakeholders.



Jeyapriya Partiban

Head of Advisory;
Head of Risk Consulting and Governance
KPMG in Bahrain

Active technology and security oversight

Principle 2

The company board actively oversees strategic AI technology choices and expects appropriate consideration of critical parameters such as speed, cost, sovereignty, privacy, safety and security.

- 2.1 Deliberate AI capability ownership
- 2.2 Well-managed vendor dependency
- 2.3 AI safety and security

Technology and software choices are not always board-level topics. However, the setup of the AI technology stack — from data centers and cloud infrastructure to the choice of foundational AI models to proprietary- or third-party data and applications — has a major impact on cost, development speed, quality and above all the control that companies have over their AI systems and solutions.

With increasingly sophisticated AI-enabled cyber threats and security risks to face, companies may wish to enhance their oversight of AI security. Boards and management need a good understanding of key processes for considering security risks from third party products and from supply chain vulnerabilities. The board needs to be briefed on, and understand, the new or changed risks associated with major technologies used by the company.

Although board members do not need to become deep technology experts, they do need an understanding of how technology choices affect the strategic position of the company. This applies in the short term (e.g. keeping up with developments and opportunities, cost and quality of the outputs) and in the long-term (e.g. dependency on large AI players). Oversight may include reviewing management’s rationale for moving quickly versus moving more cautiously, with the expectation that commercial, reputational and safety and security implications are fully considered.

Regional, and even national, sovereignty is increasingly relevant. Based on their specific geopolitical context, boards will have to factor in the relevance of sovereignty on their agenda, in particular how investment in sovereign-ready AI can open new opportunities and new risks. Boards can consider whether the company can expect governmental support or restrictions, and how this informs decisions on technology choices.

2.1

Deliberate AI capability ownership

The board expects that decisions regarding the development or procurement of AI capabilities are made deliberately and responsibly, with a clear understanding of which AI functions are strategically critical to the company’s long-term continuity, competitiveness, and resilience. These functions could be designed to use proprietary algorithms, sensitive data, or sector-specific models. Geopolitical considerations, such as microchip supply, regulatory environments, supply chain security, are also considered in the context of external or vendor, partner and supplier decisions.

2.2

Well-managed vendor dependency

The board expects management to have a thorough decision-making process to enable deliberate, informed choices about dependencies on external AI vendors, vendors using AI, and on the possible adoption of open architectures and interoperability standards. It oversees how management maintains a clear posture toward third party use of AI, and the dependencies and relationships created.

The board oversees management’s considerations of associated risks and mitigations: for example, a vendor failing or being legally prevented from supplying services, or a scenario where multiple vendors are engaged in a single project with no clear legal structure for establishing liability if the AI fails. The board expects management to regularly assess these dependencies, and to assess how they will help preserve the company’s ability to adapt to technological and market changes.

2.3

AI safety and security

The board expects management to establish risk-based and proportionate traceability, auditability (including clear audit trails), legal review and control over AI systems used to safeguard technology and data against bad actors, data poisoning, misinformation, agentic AI misalignment, privacy breaches, cyber-attacks, and the ongoing potential for negative events.

The board discusses with management the increased risk exposure that AI brings, and how employee training and business continuity planning might need to be adjusted in the light of AI-related risks. It expects that robust and resilient AI security practices are implemented. Security frameworks should address technical and systemic risks, such as operational failures, misinformation, deepfakes, sabotage, litigation risks and adverse events. Crisis management capabilities should be in place, and updated on an ongoing basis, to respond effectively to emerging threats and AI-related incidents.

Workforce transformation and human accountability

Principle 3

The company board expects that management proactively develops clear guiding principles for effective AI adoption in the organization. It considers how the AI adoption strategy affects the company's human capital resources. This includes the future role of humans vs. AI as well as the required capability building within the workforce.

- 3.1** Defined and safeguarded quality of (human) decision-making
- 3.2** Strategic workforce transformation and upskilling
- 3.3** Appropriate long-term talent pipeline

AI adoption will likely have a major impact on the workforce and culture of an organization. To achieve AI adoption at speed, companies will need to engage and train the whole workforce.

As the use of AI evolves, there will be expectations for companies to maintain constant vigilance over how to achieve balance between human judgment and AI output, and over where human judgment is non-negotiable, in alignment with the local regulatory requirements and with the company's values. This will impact the future skills requirements of the company, the understanding of where humans truly add value over and above what is delivered by AI, and of how companies develop these skills in their workforce. There will be expectations that outcomes of high impact AI decisions are explainable and understandable to employees and to other relevant stakeholders.

Although HR strategies typically fall under the remit of management, boards may need to oversee, at an appropriate strategic level, how management responds to the profound and sometimes disruptive impact AI can have on organizations and individuals.

To ensure a vibrant future labor market from which companies can recruit new talent, boards should encourage management to consider to what extent the company can contribute to the development of the talent pipeline in a broader sense and the possibilities for using AI to bridge the digital divide.

3.1

Defined and safeguarded quality of (human) decision-making

Boards ask management to define which types of decisions must always involve meaningful human oversight, or involvement, or accountability, consistent with "human-in-the-loop" and "human-on-the-loop" concepts. They challenge management to explain how the use of machine output is proportionate to the context, risk, and impact of these decisions. They consider this in the light of current and anticipated regulatory requirements and potential reputational harm. They expect that, for high impact decisions, a higher degree of AI involvement is appropriate only where outputs are explainable, auditable, and aligned with the company's values and risk appetite.

The board oversees whether AI outputs are explainable to and understandable for relevant stakeholders, especially in high-impact contexts. It expects management to oversee and report on regular assurance of decision-making using AI, including testing for fairness, transparency, and unintended consequences, and it oversees management's processes for addressing situations where risks are found. It remains informed on where innovative tools, processes and services are being developed, both "bottom up" as well as "top down" in the organization.

3.2

Strategic workforce transformation and upskilling

The Board expects management to develop a comprehensive HR transformation strategy, including considering the need for more or fewer employees overall, and by department or skill group, to support the company's long-term vision and its ability to meet strategic goals. The strategy could include plans for workforce transformation, the introduction of AI tools and agents designed to act as virtual colleagues, commitments to upskilling, reskilling, and for redeploying employees to build a sustainable, AI-augmented environment.

The board oversees how AI-driven workforce changes are managed, with a view to sustaining employee engagement, finding new ways to build and retain critical skills, protecting wellbeing, and preserving a culture of trust, responsibility, adaptability, and inclusion.

3.3

Appropriate long-term talent pipeline

The board considers the longer-term effects of AI adoption in the organization and its impact on workforce supply and demand. It oversees how management are proposing to balance trade-offs between upskilling and replacing employees, and how entry-level job descriptions and apprenticeships might evolve.

This oversight may include the impact of AI on the company's talent pipeline, within and outside the company, and on the future supply by the labor market at large. In this context, the board oversees how management assesses any appropriate training and support for people moving into new roles.

Building trustworthy AI

Principle 4

The company board expects that management proactively adopts standards for trustworthy AI that reflect the company's values and regulatory obligations, and appropriately considers the company's stakeholders and society at large.

- 4.1 Fair and inclusive AI solutions, processes and culture
- 4.2 Safeguarded privacy, data usage and protection in AI solutions
- 4.3 Impact on the physical environment
- 4.4 Trustworthy AI
- 4.5 Legally compliant AI

AI technology can have significant secondary effects for the company, as well as for its stakeholders and society in general, which might not be in line with the company's values and culture. Different jurisdictions are taking different approaches to regulating these effects and the expectations of a company's stakeholders might vary greatly depending on the markets the company chooses to operate in, meaning that companies must navigate between these different expectations.

The board plays an important role in overseeing appropriate discussion on how the company's values translate into important AI-related values such as fairness, transparency, safety, accountability, privacy and data protection. It expects management to set clear policies based on these values. In addition, the use of AI can bring with it significant sustainability impacts, including but not limited to energy, pollution, and water use. Boards should be aware of the extent of this impact and discuss with management the related risks, especially regarding long-term supply certainty and potential conflicts with third parties.

Beyond this, companies must always comply with laws and regulations in the jurisdictions in which they operate. Therefore, one of the concerns is the degree to which regulations in different countries and markets can conflict or be inconsistent, or require different approaches.

Trust is a business requirement, key to the adoption and usage of AI. A company's values-based guidelines should be reflected in the testing and oversight of actual AI systems, with a goal of lawful, ethical, robust and truly "trustworthy" AI.

Trust is a business requirement, key to the adoption and usage of AI. A company's values-based guidelines should be reflected in the testing and oversight of actual AI systems, with a goal of lawful, ethical, robust and truly "trustworthy" AI.

4.1

Fair and inclusive AI solutions, processes and culture

While enabling management to innovate at pace, the board oversees whether management has processes to check that AI systems are developed and deployed with fairness, transparency, safety, inclusivity and accountability, encouraging management to define these principles and expectations such that they reflect the company's values and strategy.

AI systems are monitored and evaluated on a risk-based and proportionate basis (including consideration of the involvement of an objective external party where appropriate) to prevent bias and discrimination, and to identify and mitigate unjust or unintended discriminatory effects.

The board fosters a culture of transparency and accountability throughout the AI lifecycle from initial design to application, including clarity on the company's external activities regarding AI-related advocacy and regulatory engagement activities.

4.2

Safeguarded privacy, data usage and protection in AI solutions

The board oversees how management establishes and maintains privacy and data usage and protection standards for the design and operation of AI systems. It oversees the organization's data and privacy governance framework, which should ensure that data is legally obtained and its use is permissioned and transparent. Additionally it oversees how data is managed throughout the value chain, including protection against unauthorised access or misuse, and the purposes for which it is managed.

The board expects management to put robust data and privacy policies in place that are aligned with legal and regulatory obligations, and with the company's standards. It checks with management that these policies include suppliers, vendors, partners and regulators, and encourages good governance practices throughout the supply chain, including vendor and distributor security undertakings and commitments and different norms for AI use in different jurisdictions.

4.3

Impact on the physical environment

The board oversees, in line with the company's commitments and its regulatory environment, how AI solutions are designed, implemented, and operated in a manner that supports the company's stated commitments to long-term sustainability and responsible stewardship. Where companies have made environmental commitments, AI initiatives are assessed in a risk-based manner for their external impacts, including resource consumption such as energy use and water impact. Where relevant, the board oversees management's consideration of the entire value chain.

4.4

Trustworthy AI

The board oversees processes to translate values-based principles into trustworthy AI. Trustworthy AI involves a combination of technical design and effective monitoring and organization for key activities such as AI learning practices, testing, output drift and incident management, resulting in systems that are lawful, ethical, and robust. AI adoption strategies in the organization should include training on responsible use of AI and clear instruction on what is not permissible.

4.5

Legally compliant AI

For companies operating across multiple jurisdictions, the board oversees how management ensures adherence to multiple, sometimes conflicting, requirements and how management adopts a clear and defensible approach.



The question is no longer whether to adopt AI, but how to translate it into real business value. The challenge is not technology - it is integration.

Organizations that will lead in Bahrain are those that connect AI to strategy, embed it into operating models, and define ownership over AI-driven decisions. When governance, data, and technology are aligned, AI becomes a driver of performance - not just experimentation.



Manav Prakash

Head of Management Consulting and
AI Advisory Services
KPMG in Bahrain

The work of the board

Principle 5

The company board actively considers how AI will affect its own board governance practices, including the board oversight process itself, and defines what effective AI oversight will look like. It adapts its governance structures and processes, especially regarding oversight of risk management, to meet confirmed AI oversight requirements.

- 5.1 Appropriate oversight capabilities for AI and other emerging technologies
- 5.2 Effective AI oversight processes
- 5.3 AI tailored risk management
- 5.4 Adherence to a globally diverse regulatory AI landscape
- 5.5 Transparent and outcome-based reporting

The introduction of AI is significantly impacting traditional board governance practices. Both board and management are entering uncharted territory where board members may be able to rely less on their experience than before. Boards need to find ways to stay up to date and to understand the impact and potential of AI, and they should expect that management does the same.

Given the speed of AI development and its deep impact on the organization, boards' traditional linear ways of working may need adjustment. In response to growing oversight complexity and urgency, boards may consider evolving their governance approach, e.g. by revising their structures and composition, including committees and their mandates, board processes such as frequency of reviews and agenda time allocation, or the use of external expertise, while retaining the value of stability and discipline in their way of working. Boards may consider using AI themselves to look at new metrics, and to support board work, providing there are clear guidelines regarding permissible use by board directors. Board committees may have to review and update their charters to specify those committees' AI-related responsibilities.

The roll-out of AI introduces new risks for the company and might also change the character of traditional risks, e.g. the speed and extent with which they will occur. Boards face the challenge of maintaining effective oversight of these risks in the context of their normal risk oversight processes, while at the same time not stifling AI-related innovation and business building.

The use and impact of AI in companies might be difficult for stakeholders to follow and understand. Boards should expect transparent, accessible, and risk-based reporting by management to the board on the use and expected impact of AI, and on how it is being governed, and appropriate communication to company stakeholders.

Given the speed of AI development and its deep impact on the organization, boards' traditional linear ways of working may need adjustment. In response to growing oversight complexity and urgency, boards may consider evolving their governance approach, e.g. by revising their structures and composition, including committees and their mandates, board processes such as frequency of reviews and agenda time allocation, or the use of external expertise, while retaining the value of stability and discipline in their way of working.

5.1

Appropriate oversight capabilities for AI and other emerging technologies

The board continuously advances its collective knowledge of AI and other emerging technologies so that board members possess the technical fluency, agility, diversity of perspective and strategic insight needed to govern AI's rapid evolution, and to seek expert advice when appropriate. It regularly reviews and, where appropriate, adapts board composition and governance structures based on the needs of the company and its strategic objectives in the emerging AI landscape.

5.2

Effective AI oversight processes

The board dynamically adapts its processes to help enable effective oversight of AI and other transformative technologies, which reflect the unique, ongoing and training oriented nature of AI projects, including the need for different types of milestones and metrics.

The board has a strategy for its own use of AI e.g. regarding information gathering and processing, and the use of AI to support its own decision-making, as well as for putting in place safeguards for issues including, but not limited to, confidentiality and legal privilege. Where appropriate, the board considers adopting a formal and legally vetted AI use policy for board members that is aligned with the company's wider AI strategy.

5.3

AI tailored risk management

The board expects that AI-specific risks are included in the risk management framework, that existing risks are reevaluated based on the impact of AI, and that potential for disruption is properly assessed. Risk management frameworks should reflect internal operational risks and external environmental, competitive, regulatory, and reputational risks from AI. They should take into account how these risks could impact the company's viability and competitiveness.

The board expects management to establish AI-specific metrics, risk metrics, operational controls, reporting lines, and escalation procedures, especially where risks are significant or outcomes are uncertain.

The board oversees management's risk mitigation strategies, including insurance, contingency planning, and incident response. These strategies may address potential data breaches, supply chain disruptions, and other AI-related threats.

5.4

Adherence to a globally diverse regulatory AI landscape

The board is briefed appropriately on current and emerging relevant regulations, and on their impact on the company in all jurisdictions where the organization and its value chain operate, including national and multi-jurisdictional AI law.

5.5

Transparent and outcome-based reporting

The board requires that management reports systematically, explicitly and transparently about where and why AI is being adopted, specifying whether it is seen as a risk, an opportunity, or both, and about how it is being responsibly integrated into strategy and risk management.

Subject to applicable disclosure regimes and materiality assessments, the board also oversees external communication of the company's AI policy and the board's own oversight of AI governance through transparent narrative reporting that clearly explains the principles applied, the shifts in business processes and risk management, and the outcomes achieved.

The board expects management to provide easily understood and accessible explanations of major AI capital allocation decisions to the board (and to company stakeholders where appropriate), including the anticipated benefits, risks, and alignment with the company's long-term strategy.

The board fosters a culture of accountability, where both successes and failures in risk-taking and capital allocation are openly discussed, and where appropriate stakeholder feedback is actively sought and considered.

Conclusion

“At the end of the day, we’re looking at a system-level shift that boards must learn to oversee”

That was the opinion of one advisor to the project, reflecting strong agreement among our advisory group that the board plays an essential role in overseeing responsible AI implementation. Our advisors also agreed that board members should approach AI with ambition, not fear.

Boards bring a broader perspective and a disciplined appreciation of how AI reshapes risk profiles, workforce expectations, and the operating environment in which companies compete. Precisely because they are not immersed in the day-to-day, board directors are uniquely positioned to see the bigger picture. They add value through measured, iterative, and multidisciplinary oversight, informed by a sensitivity to the wider ecosystem and to the long-term implications of how AI is deployed today and in the future.

These principles equip board directors with a practical decision-making framework to oversee AI with confidence amid rapid technological, regulatory, and business change. There will rarely be a single right answer to some of the questions and dilemmas that AI brings to the boardroom, but effective oversight will continue to depend on sound judgment, sustained curiosity, and active engagement with AI as a strategic and governance issue.

By anchoring boardroom discussions in these principles, board members can deliver informed, consistent, and responsible governance that positions their organizations to compete, grow, and thrive in the age of AI.



The Business School
for the World®

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

kpmg.com/bh

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.

© 2026 KPMG Fakhro, a Bahrain partnership registered with the Ministry of Industry and Commerce (MOIC), Kingdom of Bahrain 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.

© 2026 INSEAD Institut privé d'enseignement supérieur.

KPMG refers to the global organization or to one or more of the member firms of KPMG International Limited ("KPMG International"), each of which is a separate legal entity.

KPMG International Limited is a private English company limited by guarantee and does not provide services to clients. For more details about our structure please visit kpmg.com/governance.

The views and opinions of external contributors expressed herein are those of the interviewees and do not necessarily represent the views and opinions of KPMG International Limited or any KPMG member firm. KPMG's participation and contribution in this regard is not an endorsement, sponsorship or implied backing of any company's products or services.

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

Throughout this document, "we", "KPMG", "us" and "our" refers to the KPMG global organization, to KPMG International Limited ("KPMG International"), and/or to one or more of the member firms of KPMG International, each of which is a separate legal entity.

Designed by Evalueserve.

Publication name: AI Governance Principles for Boards | Publication number: 140336-G | Publication date: April 2026

