



# **The AI Governance Guide**

# Table of Contents

Introduction: Leveraging AI for  
Innovation Responsibly

**3**

Chapter 3: Assembling Your Stakeholders

**9**

Chapter 1: Understanding AI Risks

**4**

Chapter 4: Implementing AI Governance

**12**

Chapter 2: AI Governance and The Principles of  
Responsible AI

**5**

Conclusion: Maintaining AI Governance

**16**

# Introduction: Leveraging AI for Innovation Responsibly

For C-suite executives around the world, the whirlwind rise of artificial intelligence (AI) has raised the stakes on AI adoption. These business leaders rightly view AI with both enthusiasm and skepticism — yes, AI has tremendous potential to transform business, but it also carries tremendous risk.

And the greatest risk of all may be inaction, watching your competitors take advantage of this paradigm-shifting technology while getting left behind.

When innovation is a must, the question is not an “if” but a “how” — how to innovate safely and with intention. The answer? Creating the structure and guardrails to prioritize and use AI responsibly.

An AI governance program builds upon foundations within your organization, such as your existing data governance and risk management programs. It puts people at the center of your approach, equipping your team with the resources and controls to foster safe, compliant, and ethical AI use.

In the following pages, we'll explore the bedrock of principles and practical frameworks of AI governance. Learn how to tap into the power of AI without sacrificing operational integrity or stakeholder trust.

**AI bias:** Inaccurate outputs and/or outcomes resulting from an AI system

**Data hygiene:** The process of keeping data, files and databases clean, accurate, and error-free

**Data leakage:** The unauthorized transmission of data to an external recipient

**Generative AI:** AI systems that produce text, images, audio, and other outputs based on patterns in large datasets

**Hallucination:** False or misleading information presented as fact by an AI system

**Model drift:** When an AI system's performance worsens over time due to changes in data or the underlying model

# Chapter 1: Understanding AI Risks

As with the adoption of any technology, integrating AI into your organization carries a host of risks. These risks are wide in scope and multidimensional, with potential impacts that cut across business functions and departments. The first step in combatting these risks is to understand them, the ways in which they act and interact — with your organization, with third parties, and with other risks — and the threats they pose.

**Bias:** AI systems can perpetuate and amplify biases, leading to inaccurate outputs and outcomes. These biases may stem from the user, underlying algorithms, or problems within your training data.



**Cybersecurity:** The adoption of new AI tools and platforms increases an organization's attack surface (the vectors through which a threat actor may seek to access or impact your systems). Many threat actors are also tapping into the power of AI to execute increasingly sophisticated cyberattacks.



**Environmental:** AI consumes significant energy resources — not only electricity, but also water to cool data centers and manufacture semiconductors. Balancing innovation with sustainability remains a challenge, particularly for organizations focused on reducing their negative environmental impact.



**Financial:** Developing, maintaining, and using AI systems can be extremely expensive, leaving companies with a potentially unexpected and unsustainable financial burden.



**Hallucinations:** AI may present false or misleading, yet plausible information as fact, resulting in erroneous outputs or outright fabrication. Hallucinations typically result when the AI cannot draw on a strong factual foundation.



**Job Displacement:** AI — generative AI and AI agents, in particular — has driven widespread fears of job displacement. In this evolving environment, many workers will need upskilling and reskilling as AI takes on an increasingly prominent role in the workplace.



**Legal and Regulatory:** As AI regulations proliferate around the globe, legal and compliance teams may struggle to stay on top of requirements. The same principle applies at the ground level — every employee needs to understand their legal responsibilities and how to remain compliant when using AI as part of their job.



**Misinformation:** AI can be used to rapidly create and spread false information, eroding trust and making it difficult to identify accurate information.



**Privacy:** Improper collection, management, and/or use of personal data can infringe on and even violate individual privacy rights. These violations may be driven by unclear usage policies, lack of consent mechanisms, the deanonymization of sensitive information, data breaches, and/or data leakages.



**Access to Data:** In addition to incorporating data privacy requirements, overall access to larger data elements must be considered as AI is used in day-to-day operations and tasks that connect data from multiple systems. The access controls associated with roles and responsibilities at individual data levels need to be incorporated as data is connected or rolled up for use in AI solutions.

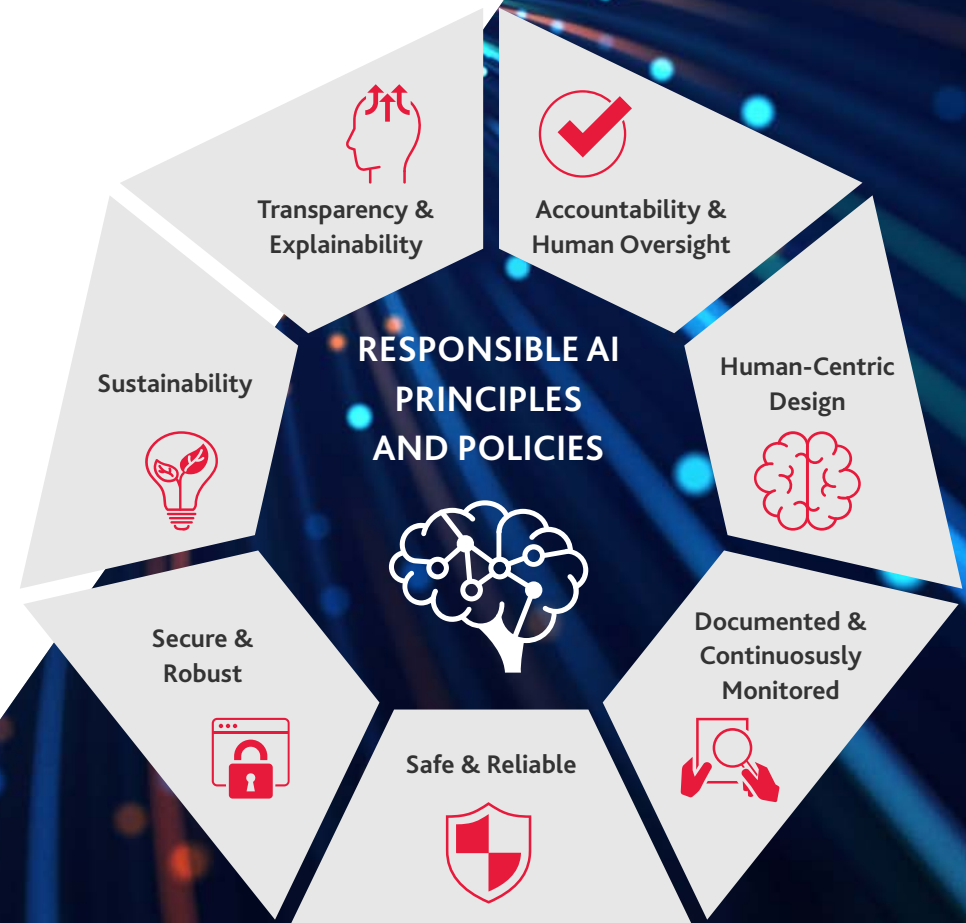


# Chapter 2: AI Governance and The Principles of Responsible AI

At the foundation of your approach to responsible AI are the core principles that define what responsible use looks like. These principles underpin the governance and guardrails embedded within your technology and processes, as well as the resources and training you provide to your employees. They help you assess and mitigate risk, implement effective controls, and assist with regulatory compliance.

In addition to content covered in this article, companies may also consider utilizing the NIST AI Risk Management Framework and ISO 42001. ISO 42001 is an international standard that specifies the requirements for establishing, implementing, maintaining, and continually enhancing an Artificial Intelligence Management System (AIMS) within an organization. These resources provide valuable frameworks for the responsible development and use of AI systems by effectively managing the risks associated with an organization's AI applications.

BDO's Responsible AI Principles help guide at any stage of your adoption journey, from design and development through implementation and ongoing maintenance. We not only help our clients adopt these principles — we are aligned with this framework at BDO for our own AI initiatives.



## TRANSPARENCY AND EXPLAINABILITY

Transparency and explainability build trust and foster responsible AI use. A transparent and explainable AI system has a clearly defined purpose, function, capabilities, limitations, and impact. Its methods, data, assumptions, and decision-making criteria are extensively mapped out.

Organizations should capture this information with extensive documentation, covering not only what the AI system does, but also why and how it does it. These documents must be disseminated to all relevant stakeholders, including technical teams, business stakeholders, and end-users (see: Documented and Continuously Monitored).



The principles of transparency and explainability are especially important to end-users' interactions with AI systems. Whether these users are engaging with an AI system or AI-generated content, such as text, images, audio, or video content, they must be informed about the extent to which their experience involves AI, system limitations, potential margins of error, and how their information will be used and stored.



## ACCOUNTABILITY AND HUMAN OVERSIGHT

Clear accountability and responsibility for AI systems — not simply their use, but also their design and deployment — are key to maintaining proper oversight of AI-supported processes and decision-making. While AI can yield powerful tools and more efficient business processes, the technology may pose significant risk without proper oversight and governance.

Human oversight helps ensure that AI systems align with organizational values and ethical considerations. Establishing clear lines of responsibility and accountability can act as a safeguard against potential misuse or unintended consequences.

## HUMAN-CENTRIC DESIGN

The rise of AI presents a unique opportunity to make this powerful technology as beneficial for as many people as possible. As organizations figure out how to responsibly implement AI, they must not lose sight of their obligation to put people at the center of their approach.

This means building systems that support and provide the training and upskilling necessary to incorporate AI into day-to-day work. A human-centric design also includes privacy policies that protect personal information and ensure people maintain ownership over their own data.

## DOCUMENTED AND CONTINUOUSLY MONITORED

A written record of AI systems enables an organization to clearly identify, communicate and report on where and how they are using AI. It should include:

- ▶ Development process, protocols, and lifecycle governance
- ▶ AI models being used
- ▶ Intended use
- ▶ Stakeholders, roles, and responsibilities
- ▶ Design decisions
- ▶ Training, validation, and testing datasets used

This record should not be static. Rather, it must be regularly updated to reflect relevant events and modifications that occur as the AI system and its use evolve. Stakeholders must continuously monitor, document and report on the system's performance, outcomes, improvements, and impact. Documentation and reporting channels are also key for maintaining compliance with applicable legal, ethical, and business standards.

## SAFE AND RELIABLE

To be safe and reliable, AI systems must not only perform consistently and accurately under normal and foreseeable conditions, but also under anomalous circumstances. Continuous testing can validate that an AI system is functioning as intended. By defining potential failures of their AI systems, organizations can prepare mitigation measures and contingency plans to overcome these potential risks. This includes accounting for unintended manipulation and misuse.



## SECURE AND ROBUST

Secure and robust AI systems can withstand and recover from threats to their integrity, functionality, or availability. Such systems must also be compatible with an organization's existing systems and processes to avoid the possibility that AI may compromise, disrupt, or degrade the organization's other tools, resources, and offerings.

"Security by design" is an engineering best practice that incorporates security from the outset of the development process, rather than treating it as an add-on or afterthought. This proactive approach not only promotes robust security design, but also the adoption of security strategies and tactics to help mitigate risks and overcome attacks.

By embedding security considerations into every phase of development, organizations can identify and address vulnerabilities early, reducing both the cost and complexity of security implementations. This approach is particularly crucial for AI systems, where compromised data and biased outputs could damage vital processes, reputation, and user trust.

## SUSTAINABILITY

The integration of AI into organizations presents a variety of considerations, particularly in the context of sustainability. For instance, as AI usage grows, it may lead to increased energy consumption, which could impact availability of resources. Organizations might need to consider alternative energy resource options, such as solar, wind, and battery storage, to support their AI operations.

Additionally, the expansion of AI and the development of new data centers can have implications for water resources globally. Data centers often require substantial water for cooling purposes. With the challenges posed by climate change and rising demand, there is an opportunity to explore innovative, sustainable, and recyclable solutions to manage water use effectively.

AI offers transformative potential and the Chief Sustainability Officer can play a crucial role in this decision-making process by providing critical insights into the environmental footprint of AI implementations, helping their organization balance technological advancement and costs with environmental responsibility.



# Chapter 3: Assembling Your Stakeholders

AI governance is not owned by any single team. It is an interdisciplinary effort, drawing on business leaders with a range of knowledge and experience. By mapping AI needs and impacts across an organization, these leaders can make the right investments to support their broader business goals. This means aligning IT, financial, cybersecurity, and other business operations to help ensure they drive toward the same vision of success.

## **Audit and Assessments:**

Ensure AI systems comply with internal and external standards through regular evaluations and assessments.

## **Cybersecurity:**

Protect AI systems from cyber threats and vulnerabilities, preserving data integrity and security.

## **Enterprise Risk**

**Management:** Convene teams and stakeholders to track, prepare for, and mitigate AI risk holistically across the organization.

## **Information Technology:**

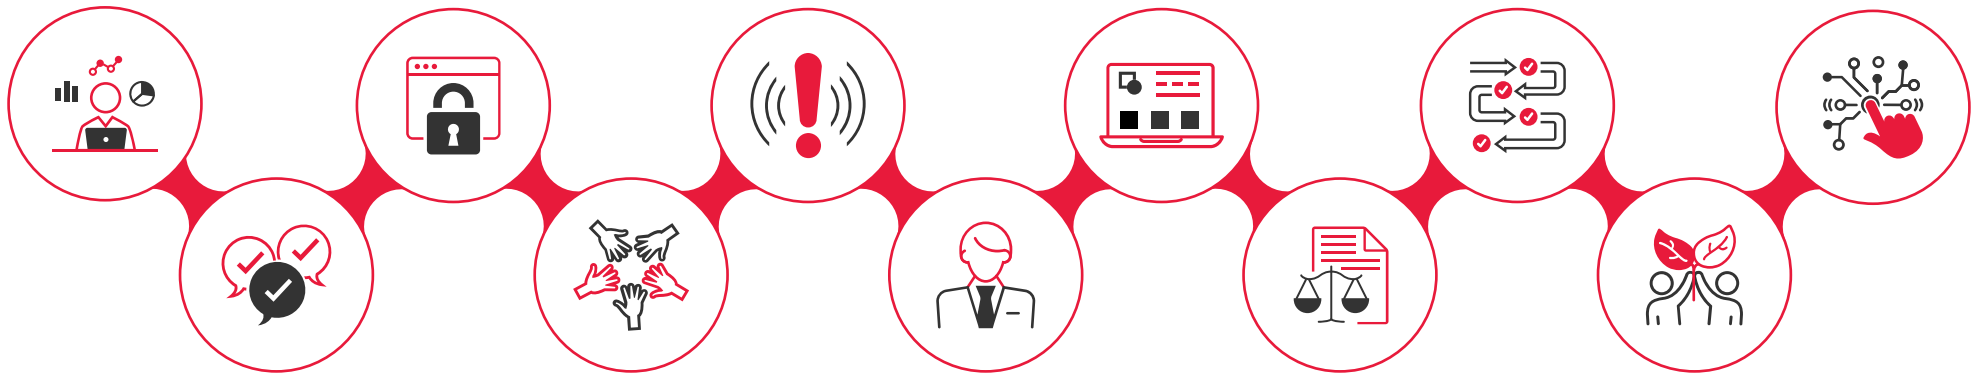
Support the technical infrastructure to help ensure AI systems are reliable, scalable, and aligned with standards and requirements.

## **Operational:**

Integrate AI into business processes efficiently and safely, while equipping employees with the necessary resources to maintain responsible AI standards.

## **Third-Party:**

Manage relationships with AI vendors and partners; validate responsible AI practices throughout the third-party ecosystem.



## **Communications:**

Disseminate information regarding AI initiatives; promote transparency and foster trust through clear and consistent messaging.

## **Diversity, Equity, and**

**Inclusion (DEI):** Evaluate the impact and potential AI bias toward certain groups of individuals, often by performing AI risk assessments.

## **Human Resources:**

Facilitate user adoption by addressing workforce impacts, providing training, and promoting a culture of responsible AI use.

## **Legal:**

Navigate the increasingly complex regulatory landscape, while maintaining organizational compliance with laws and ethical guidelines to mitigate legal risk.

## **Sustainability:**

Promote environmentally safe development goals and energy-efficient practices.

# Data and AI Strategy Committee

For companies that have one, the data and AI strategy committee plays a pivotal role in shaping an organization's future by setting a visionary AI strategy. This committee should diligently survey employees and customers to gain critical insights into current and future AI applications, helping to ensure the organization remains at the forefront of innovation. By deeply understanding competitors and market demands, this committee positions its organization to adapt in an ever-evolving landscape. While larger companies may already have leaders with the experience and knowledge to serve in these roles, smaller or less mature companies may consider hiring people with relevant capabilities.

This committee should include leaders from across the organization who work together to define the scope and purpose of AI efforts, determine the right governance structures, and address critical questions as they arise. In larger organizations, there may be dedicated C-suite leaders in these roles, and in smaller organizations there may be management-level team members. The data and AI strategy committee typically provides updates to the organization's board of directors on a quarterly basis. The committee's role is to advise and recommend action. It may have its own decision-making authority or may defer to the organization's board of directors to make final decisions.

## EXAMPLES OF C-SUITE POSITIONS EMERGING IN LARGER ORGANIZATIONS:

- ▶ Chief Data Officer (CDO)
- ▶ Chief AI Officer (CAIO)
- ▶ Chief Privacy Officer
- ▶ Chief Strategy or Revenue Officer
- ▶ Chief Information Officer (CIO)
- ▶ Chief Technology Officer (CTO)
- ▶ General Counsel (GC) or Chief Legal Officer (CLO)
- ▶ Chief Trust Officer (CTrO)
- ▶ Chief Sustainability Officer (CSO)
- ▶ Business Line Leaders

While the CAIO is typically responsible for setting AI strategy, the Chief Privacy Officer is often responsible for managing AI governance, risk and compliance. In larger organizations, this responsibility might also lie with a team responsible for not just privacy, but also AI, cybersecurity and data governance as well.



## INTERNAL ADVOCATES

Any AI adoption effort must take a user-first, employee-centric approach to gain buy-in. If your employees don't understand your AI strategy and goals, they'll invent a narrative of their own.

Internal advocates at multiple levels throughout an organization can articulate the AI adoption plan to their teams, emphasizing the value proposition for their specific roles. Change management teams can develop these internal advocates by training them on AI capabilities and giving them hands-on experience with relevant tools. By articulating how AI will support teams, these internal advocates will encourage adoption and help assuage employees' fears.

### SPOTLIGHT: GENERAL COUNSEL

A strong partnership with your legal department is key to avoiding regulatory or geographic compliance related issues throughout the adoption process. Getting their blessing on your approach, rather than asking for approval on what you've done, will prevent disagreements and costly disruptions further down the road.

Bringing in the legal department early will have the added benefit of turning their team into de facto technical subject-matter experts (SMEs). As the use cases and regulations surrounding AI become more complex, this technical acumen will become increasingly vital. This is one reason why some organizations are appointing a Chief AI, Privacy, Cyber & Data Governance Leader who reports directly to the General Counsel.

### SPOTLIGHT: BOARD OF DIRECTORS

Management is responsible for creating the appropriate control environment to mitigate risks. To provide proper oversight, boards must become well-versed in the fundamentals of AI. They must possess a foundational understanding of the technical components of AI systems and their broader implications for the organization. With this knowledge, boards can facilitate progress by providing oversight and guidance to innovate successfully.

Boards should have access to the appropriate resources – external and internal – to help them assess the organizations' AI strategies. What is the purpose of our AI investment? Is it to enhance operations? Drive revenue? Boost customer loyalty? Understanding these specific goals will help boards determine whether AI strategy aligns with overall corporate objectives.

To determine whether the right foundations are in place to support AI initiatives, boards should assess the organization's infrastructure, talent, and resources to ensure they are equipped to handle the demands of AI deployment. These questions will help ensure an organization is prepared to not only implement AI but also to drive innovation and create value.

# Chapter 4: Implementing AI Governance

When creating a comprehensive AI governance framework, organizations should build on strong and reliable existing foundations, policies, and procedures. This framework will help you align your AI investments with broader business goals, while maintaining a focus on trust, strategic objectives, and ethical use. This framework should also enhance resilience, helping your organization adapt to changes in the AI landscape and emerging laws and regulations.

Beyond providing adoption and operational guidance, demonstrating the existence of a strong framework will also help build trust and confidence among stakeholders, including employees, clients, partners, and regulators. This trust is crucial for fostering an environment where innovation can thrive.

## RISK MANAGEMENT

Effective risk management is a cornerstone of responsible AI adoption. By implementing a comprehensive risk management plan, organizations can proactively identify and mitigate potential negative impacts on their operations, safety, reputation, and more. Risk management can also help organizations address ethical concerns, maintain regulatory compliance, and build stakeholder trust.

The first step is performing a risk assessment to identify potential risks and develop controls to mitigate them. By embedding robust risk management activities into the AI lifecycle, organizations can safeguard against potential pitfalls and ensure the responsible deployment of AI systems.

Implementing a continuous monitoring system enables organizations to track risk mitigation efforts and respond proactively to emerging challenges. This approach enables companies to detect anomalies and potential threats earlier, allowing for swift intervention before these issues escalate into more significant problems.

## DATA QUALITY AND HYGIENE

Sound data governance from the outset is vital for the successful and responsible adoption of AI. An AI system based on an underlying foundation of inaccurate or corrupted data may produce outputs that disrupt business operations, hamper decision-making, and harm users, businesses, and the public.

Understanding the inventory of your data assets and establishing quality standards, policies, and guidelines for data management is essential. Organizations must regularly validate, clean, and standardize their data, as well as perform ongoing maintenance and quality monitoring.

Implementing these standards requires the right tools and processes, as well as comprehensive training for personnel. Continuously monitoring data quality can help ensure that data used by AI systems meets the necessary standards and supports the overall success of AI initiatives.

### Assessing Data Quality



## DATA PROTECTION AND DATA SECURITY

Protecting AI systems and data from cyber threats and unauthorized access is key to maintaining data privacy and information security. Comprehensive privacy and security policies, underpinned by strong security controls such as encryption, access controls, and data anonymization, can help organizations maintain clarity and accountability regarding how their systems and data are used.

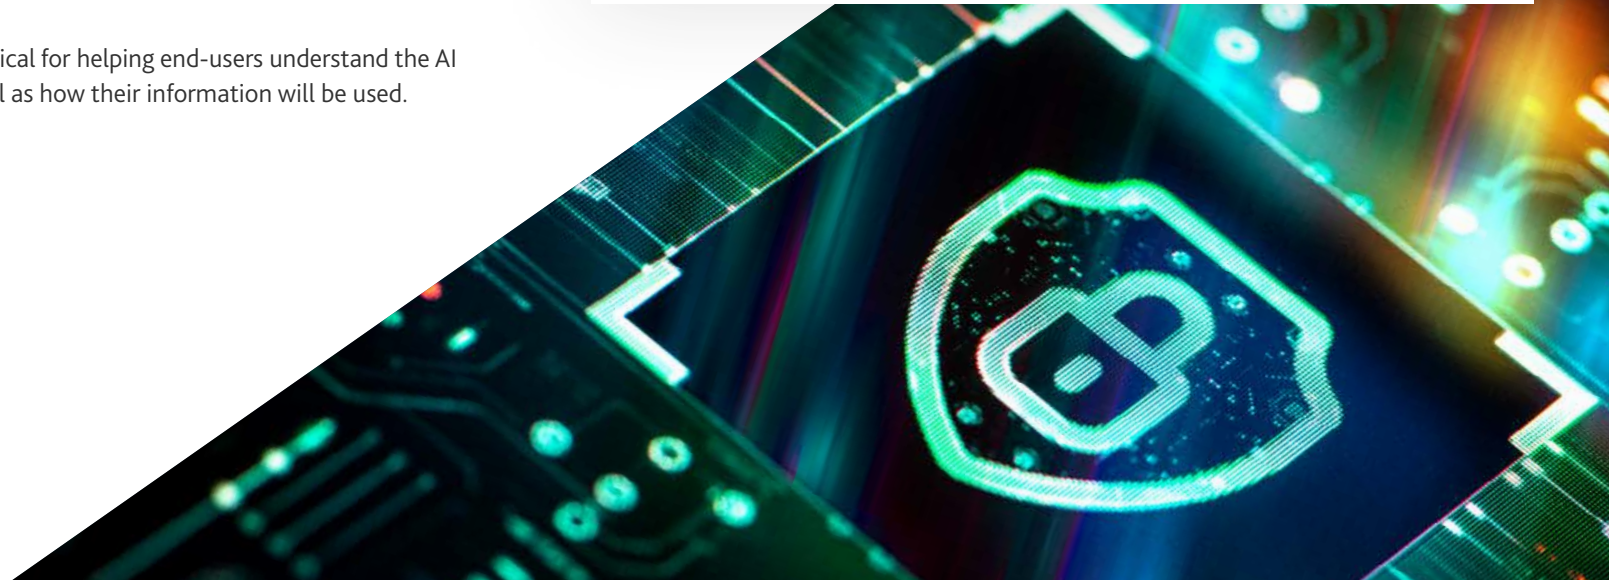
AI models are generally trained using large volumes of data. User prompts from interacting with the model are also stored and may be used to improve the model over time. Generative AI solutions also create new data in the form of the outputs they produce. Protecting data and having transparency and appropriate consent for its use are critical priorities for preserving the value of these solutions and remaining compliant with regulations. Organizations must implement security controls such as role-based access, regular access reviews and data loss prevention so that the data is protected while at rest and in transit. In addition, guidelines are needed for deciding when it is and isn't appropriate to use certain data as part of an AI solution.

Clear guidelines for data collection and use, disseminated to employees through effective, comprehensive training, will promote responsible AI practices. Monitoring systems and reporting mechanisms for ethical concerns or misuse can help facilitate compliance with internal policies and external regulations. Incident response plans and regular exercises can prepare employees to act swiftly and effectively in the event of a security breach.

Informed consent guidelines are also critical for helping end-users understand the AI systems with which they interact, as well as how their information will be used.

## 7 PRINCIPLES OF PRIVACY BY DESIGN

- ▶ **Proactive not Reactive; Preventative not Remedial:** Anticipate and prevent data privacy incidents and address privacy risks before they materialize.
- ▶ **Privacy as the Default:** Privacy is built into the system by default, requiring no action on the part of the individual to protect their privacy.
- ▶ **Privacy Embedded in Design:** Embed privacy in the design and architecture of IT systems and business processes as a core functionality.
- ▶ **Full Functionality:** Positive-Sum, not Zero-Sum: Privacy protections should not and do not need to come at the expense of security or functionality.
- ▶ **End-to-End Security & Lifecycle Protection:** Embed strong security measures throughout the information management lifecycle, from cradle to grave.
- ▶ **Visibility and Transparency:** Provide assurance to all stakeholders — users and providers — that data is being used in accordance with stated principles and objectives, subject to independent verification via a compliance and redress mechanism.
- ▶ **Respect for User Privacy:** Take a user-centric approach to data privacy, prioritizing individual privacy interests and communicating effectively.



## EDUCATION AND AWARENESS

Education and awareness initiatives can prepare employees to integrate AI into their work responsibly and effectively. These efforts may involve both general and role-specific training on using AI, understanding policies and guidelines, and meeting compliance and reporting requirements. In addition to formal education and training, successful integration requires embedding AI into an organization's culture and processes through cross-functional buy-in and a collaborative approach.

Organizations should apply specialized knowledge — spanning domains such as technology, ethics, and business processes — to develop a comprehensive educational campaign around the adoption of responsible AI. While the human resources department may take the lead on developing these trainings, they must draw on knowledge and experience from many different teams and perspectives.

Once developed, educational initiatives must be rolled out to reach as many people across the organization as possible. Regular updates will likely be necessary to both keep pace with evolving technology and incorporate feedback.

### Sample AI Curriculum



#### Introduction to AI

These courses are designed to ensure a basic foundational understanding for employees to learn the basics of AI. They outline the basics of AI technology, potential applications, and provide knowledge to be conversant on the topic.



#### AI Governance

These courses help develop an understanding of the best practices for planning, guiding, and maintaining AI systems and instilling accountability for AI systems and technologies.



#### Privacy, Data Protection, and Security

These courses introduce privacy concerns regarding the responsible use of AI and data protection education for secure development and deployment of AI systems.



#### AI Ethics

These courses are designed to give the learner an understanding of ethical principles to consider when developing and/or using AI solutions.



#### AI Bias

These courses are designed to provide education on different types of bias, how they can be introduced into AI systems and how users of systems can ensure that fairness is incorporated when utilizing AI in business decisions or analysis.



#### AI Transparency and Explainability

These courses are designed to educate on the importance of transparency when designing and deploying AI systems. They also include content related to challenges of explaining why an AI system produced a particular output.

## SYSTEMS DEVELOPMENT AND TESTING

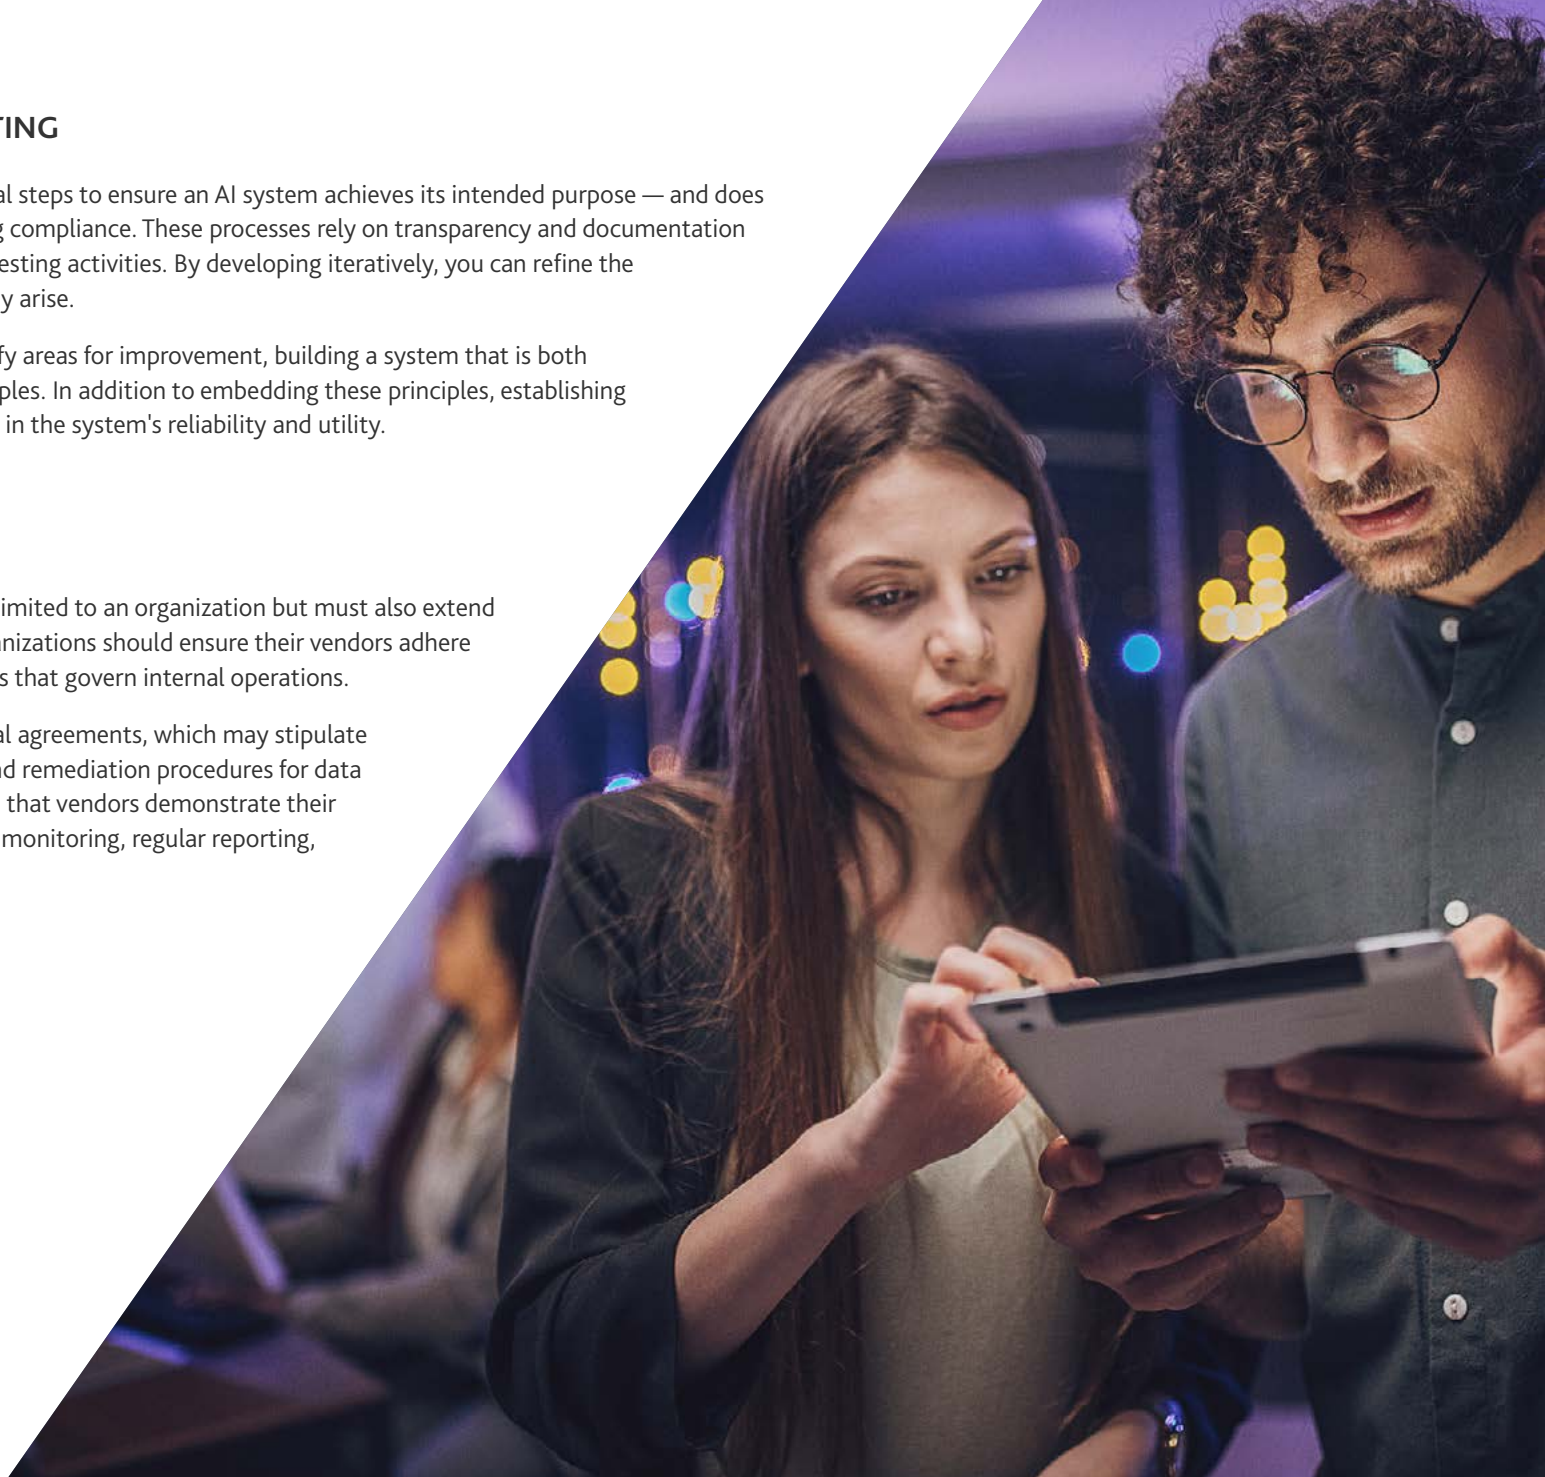
Development, testing, and validation are crucial steps to ensure an AI system achieves its intended purpose — and does so safely, responsibly, and without jeopardizing compliance. These processes rely on transparency and documentation to provide a clear record of development and testing activities. By developing iteratively, you can refine the technology as needed, addressing issues as they arise.

Rigorous testing helps mitigate bias and identify areas for improvement, building a system that is both effective and aligned with responsible AI principles. In addition to embedding these principles, establishing performance requirements can help build trust in the system's reliability and utility.

## THIRD-PARTY MANAGEMENT

The principles of AI governance should not be limited to an organization but must also extend to third-party vendors' systems and data. Organizations should ensure their vendors adhere to the same data quality and hygiene standards that govern internal operations.

These standards can be written into contractual agreements, which may stipulate quality metrics, acceptable threshold levels, and remediation procedures for data quality issues. Organizations can also mandate that vendors demonstrate their ability to maintain quality through continuous monitoring, regular reporting, and periodic assessments.

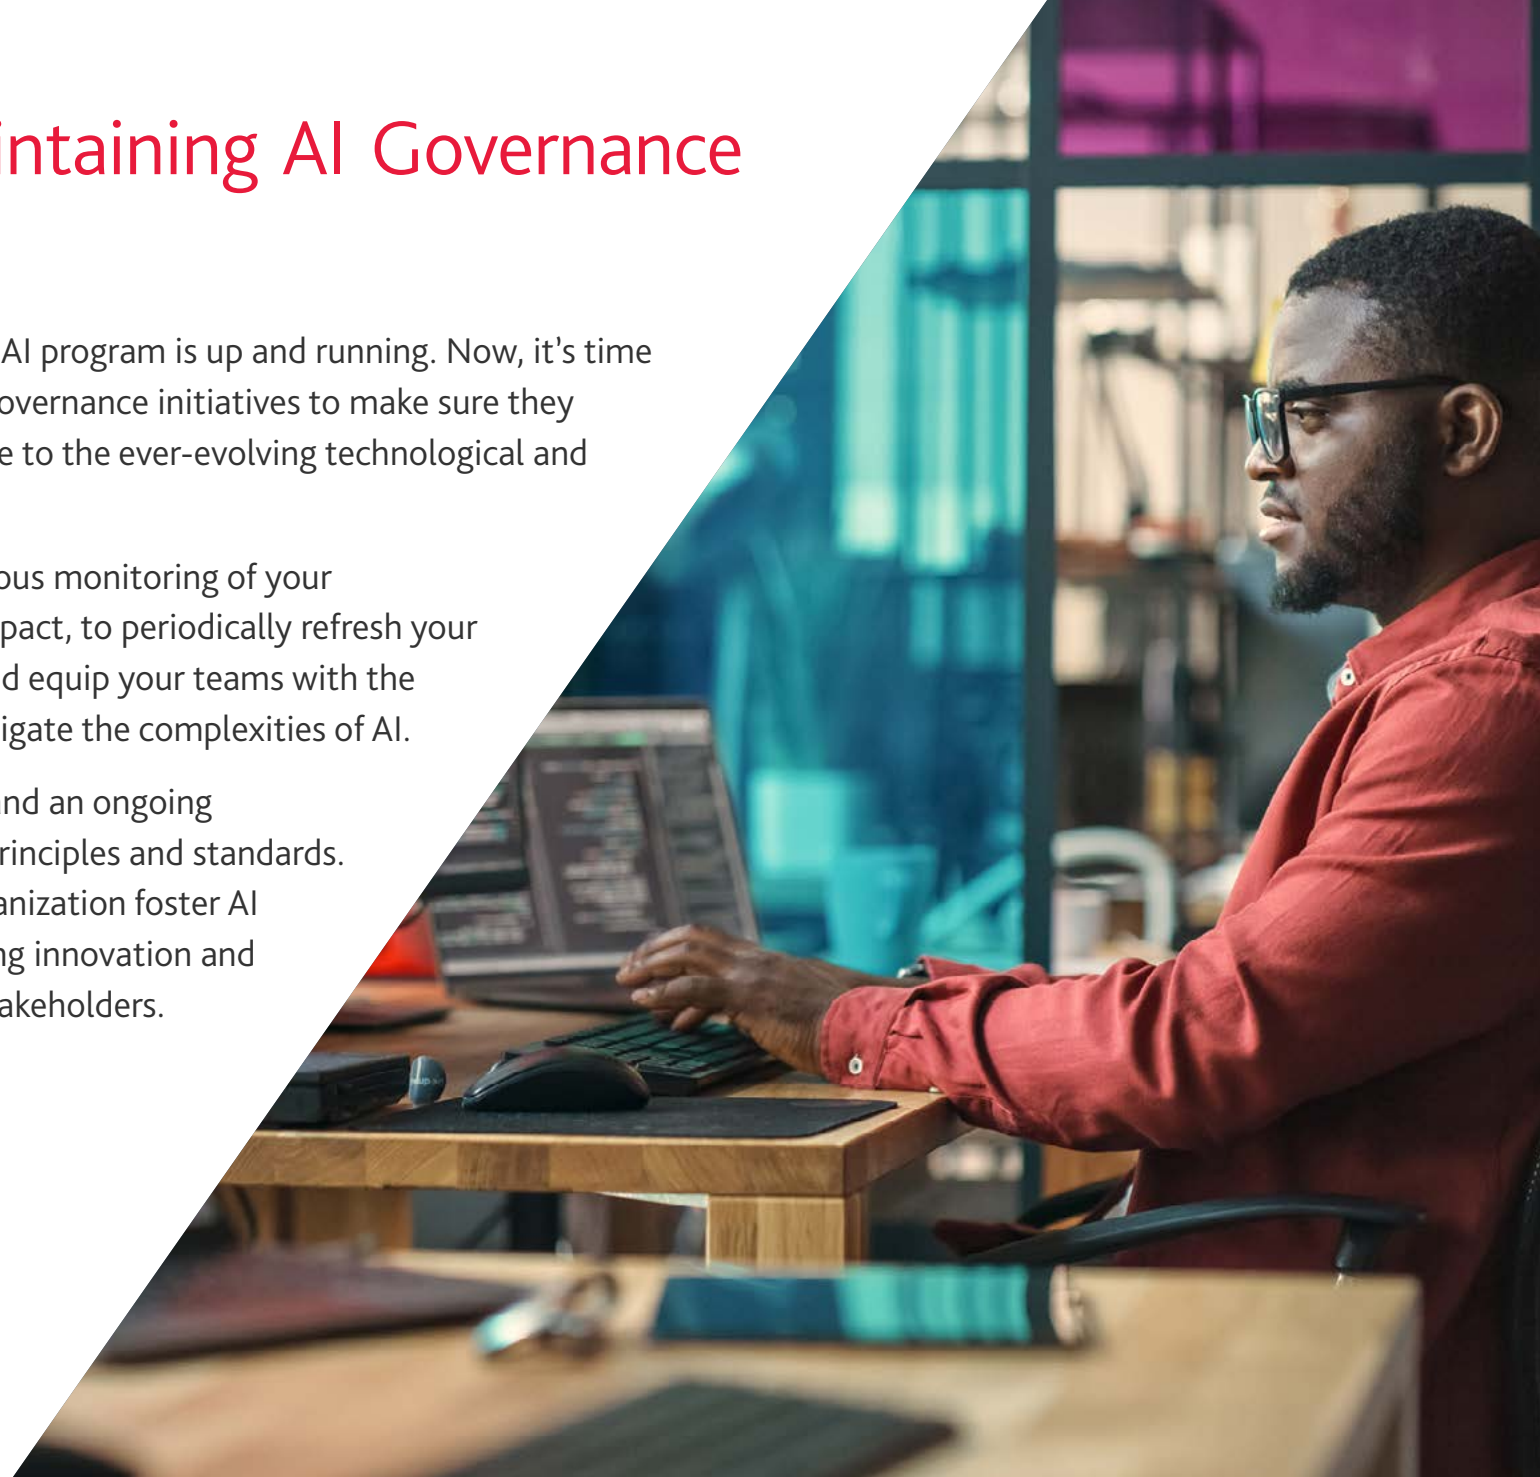



# Conclusion: Maintaining AI Governance

The work doesn't stop when your AI program is up and running. Now, it's time to maintain it: revisiting your AI governance initiatives to make sure they stay current, useful, and adaptable to the ever-evolving technological and regulatory landscape.

This means acting on the continuous monitoring of your technology's performance and impact, to periodically refresh your policies and training programs, and equip your teams with the latest knowledge and skills to navigate the complexities of AI.

AI governance requires vigilance and an ongoing commitment to upholding your principles and standards. This dedication will help your organization foster AI as a force for good business, driving innovation and creating long-term value for all stakeholders.





Our purpose is helping people thrive, every day. Together, we are focused on delivering exceptional and sustainable outcomes and value for our people, our clients and our communities. BDO is proud to be an ESOP company, reflecting a culture that puts people first. BDO professionals provide assurance, tax and advisory services for a diverse range of clients across the U.S. and in over 160 countries through our global organization.

BDO is the brand name for the BDO network and for each of the BDO Member Firms. BDO USA, P.C., a Virginia professional corporation, is the U.S. member of BDO International Limited, a UK company limited by guarantee, and forms part of the international BDO network of independent member firms. For more information, please visit: [www.bdo.com](http://www.bdo.com).

© 2025 BDO USA, P.C. All rights reserved.

