



# The AI Inflection Point

How to adopt AI responsibly  
in your organisation.



# Contents

<b>I. Introduction: The imperative for adopting AI responsibly today.....</b>	<b>3</b>
<b>II. Framework overview: Building a scalable, ethical AI future.....</b>	<b>4</b>
1. Assess: Organisational readiness and selecting responsibly built AI technology.....	5
1.1 Evaluate organisational readiness.....	5
1.2 Select AI tech that is built responsibly.....	6
2. Pilot: Identifying and piloting high-impact use cases.....	8
2.1 Identify and prepare priority use cases.....	8
2.2 Pilot against business and responsible AI criteria.....	8
3. Adopt: Integrating AI responsibly across the organisation.....	9
3.1 Train and enable the organisation.....	9
3.2 Deploy with responsibility in mind.....	10
4. Monitor: Continuous oversight and improvement.....	11
4.1 Monitor performance against business and responsible AI benchmarks.....	11
4.2 Ongoing risk management.....	12
<b>III. Embedding best practices in your organisation.....</b>	<b>16</b>
1. Employee use guidance: .....	16
1.1 Data sensitivity: .....	16
1.2 Transparency in AI usage: .....	16
1.3 Account management policies: .....	16
2. Vendor evaluation: Sample questions.....	16
3. AI governance levers .....	18
<b>IV. Responsible implementation builds responsible innovation.....</b>	<b>19</b>

Return to this page by  
clicking this menu icon



## I. Introduction: The imperative for adopting AI responsibly today.

As AI becomes an industry-wide catalyst for transformation, executives face unprecedented pressure to innovate rapidly, address competitive threats and drive operational efficiencies. However, the race to adopt AI within enterprises introduces new risks. Without careful oversight, this rapid AI implementation can lead to regulatory missteps, operational disruptions and long-term reputational damage. Balancing the drive for speed with the imperative of responsibility is no longer a trade-off – it is a strategic imperative.

Governing AI and managing its risks often feels like a herculean task. With new guidelines, frameworks and policies; and evolving international, federal and local legislation; these complexities can often obscure where organisations should begin and create challenges for stakeholders from the outset. At Adobe, our experience in responsible innovation, grounded in our [AI Ethics principles](#) of accountability, responsibility and transparency, have provided us with deep insights into navigating these challenges.

Our experience has shown that while the road to responsible AI innovation may appear daunting, success is within reach with the right tools, strategy and mindset.

One of the core decisions organisations face as they develop their AI strategies is whether to build, buy or customise an AI solution, or a combination of all three. The following approach hones in on organisations aiming to buy AI solutions and aims to build on existing values and business practices for organisations seeking to source AI solutions externally – meeting stakeholders where they are today. Grounded in independent research and informed by expert interviews on AI governance, the framework offers an actionable path forward, helping organisations assess their current standing and provide best practices for embedding responsible AI principles across the enterprise. It includes practical steps for establishing employee generative AI use guidelines, evaluating vendors through robust questionnaires and updating governance processes for AI to keep pace with the evolving landscape.

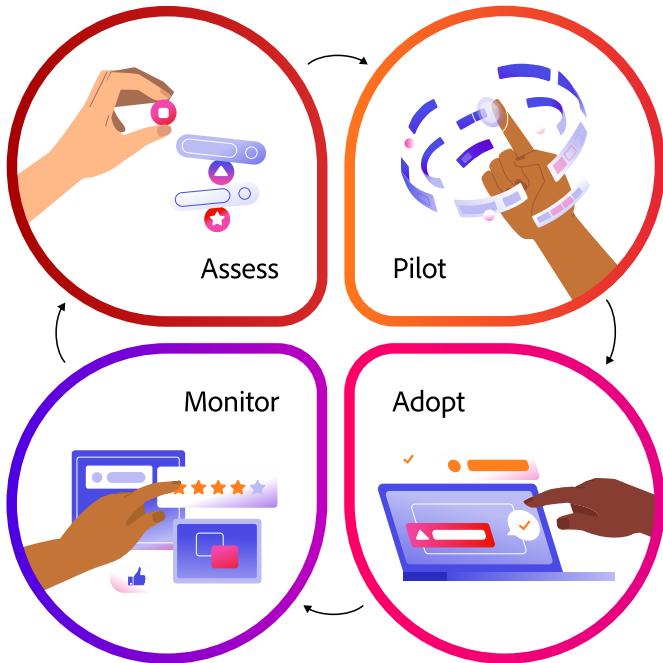
No matter where you are on your AI journey – whether assessing your organisation's AI readiness or refining existing strategies – this framework provides a proven approach, blending human ingenuity and cutting-edge AI governance to scale responsibly. By following this roadmap, organisations can assess, pilot, adopt and monitor AI solutions effectively, building a resilient foundation that fosters trust, mitigates risk and drives sustained business value.

## II. Framework overview: Building a scalable, ethical AI future.

Successful generative AI implementation requires more than a checklist of actions – it requires a strategic, layered approach where each phase builds on the last, creating a foundation for sustainable innovation and ethical AI practices. This framework serves as a series of interlocking building blocks, designed to integrate responsible AI practices at every stage – from assessing organisational readiness to scaling effectively and continuously monitoring AI systems.

Rather than seeing AI adoption as a process-driven exercise, this framework focuses on building systems that evolve in harmony with your organisational needs. It emphasises the critical balance between human oversight and advanced AI technology, ensuring that organisations can leverage AI's potential, while also aligning with ethical, regulatory and operational goals.

Each phase within this framework – readiness assessment, responsible piloting, scaling adoption and ongoing monitoring – supports long-term success, as integrated pillars that reinforce one another at every step. By embedding responsible AI practices at each phase, companies can navigate the complexities of AI adoption while fostering trust, transparency and accountability.



### Grounded in expertise and informed by research.

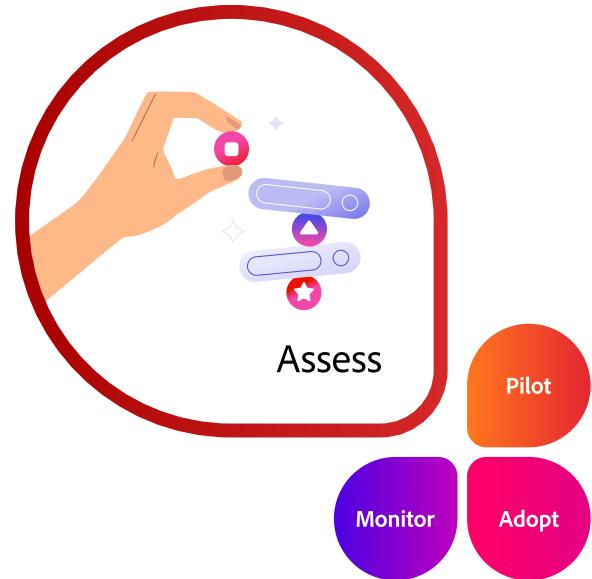
Adobe engaged an independent research firm to survey generative AI adoption, gathering insights from over 200 IT, organisational and compliance leaders across diverse industries. The research highlights current practices, challenges and successful strategies in AI adoption. Additionally, Adobe conducted in-depth interviews with industry experts and reviewed global standards, including the [European Union AI Act](#), [NIST AI Risk Management Framework](#), [Singapore's AI Verify](#), [IEEE Standard 7000](#) and [ISO 42001](#). These efforts ensure that the framework is applicable across industries and organisational sizes, regardless of AI adoption progress.

## 1. Assess: Organisational readiness and selecting responsibly built AI technology.

The journey to adopting AI responsibly begins with the people who will lead it. The **assess phase** empowers decision-makers with the tools, data and insights needed to evaluate how AI fits into your strategic priorities. This phase enables cross-functional leaders to examine the organisation's technical infrastructure, governance frameworks and AI literacy, which helps determine overall readiness.

### 1.1 Evaluate organisational readiness.

While many organisations have initiated AI adoption, of those surveyed only 21% have fully developed their responsible AI priorities, with 78% still in progress or in the planning stages, underscoring a clear need for readiness approaches. Leaders in IT, compliance, risk management and strategy are essential to build a foundation in responsible AI. This begins with a comprehensive review of the organisation's governance frameworks and AI literacy to identify gaps that could impact AI adoption.



Organisations need to take a **holistic approach** to evaluating AI readiness, blending **top-down leadership initiatives** with **bottom-up feedback** from employees who engage with AI daily.

#### Actions for readiness:

**Conduct a comprehensive readiness audit** – Evaluate the organisation's technical infrastructure, governance standards, AI-related policies, responsible innovation frameworks and compliance practices to identify strengths and areas for improvement – ensuring alignment with both strategic goals and the demands of adopting AI responsibly.

**Identify and address key gaps collaboratively** – Document additional AI policy needs in security, privacy, legal, compliance and transparency standards while engaging cross-functional teams – including IT, legal, compliance and business units – prioritise actionable next steps.

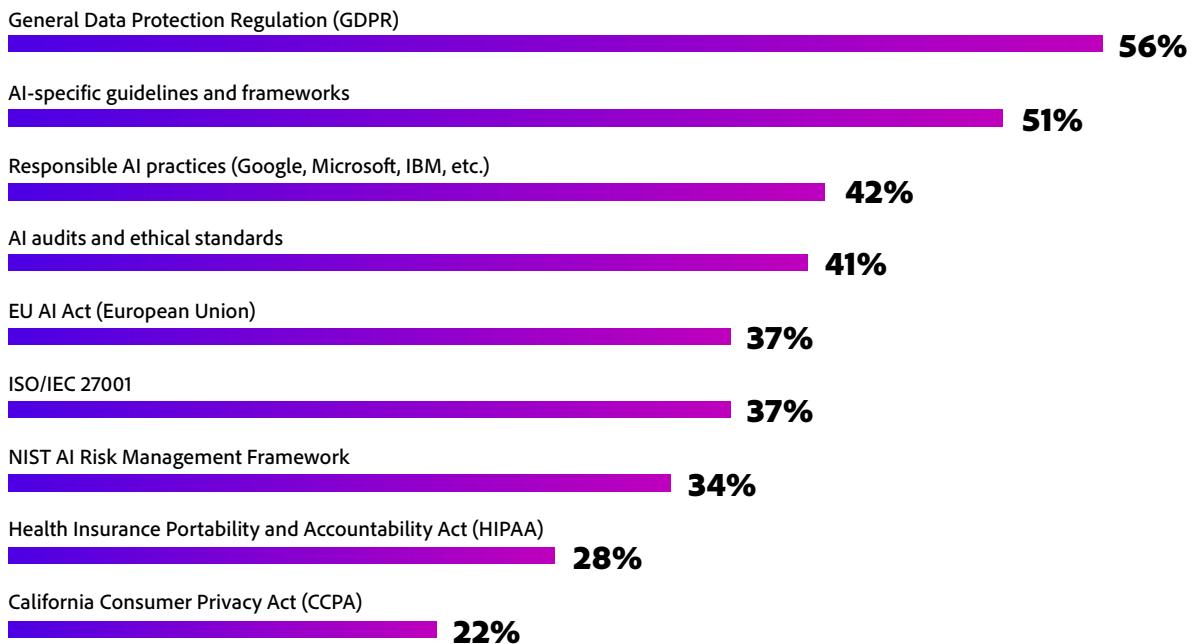
**Establish and empower governance teams** – Designate teams to oversee AI governance, ensuring compliance with both internal responsible AI standards and external regulatory frameworks – equipping these teams with the authority and resources to proactively manage risks and adapt to evolving requirements.

## 1.2 Select AI tech that is built responsibly.

Begin with a thorough review of your company's existing governance standards. These standards likely already encompass key areas such as **privacy, security, accessibility** and **legal considerations**. Global **benchmarks** like the **General Data Protection Regulation (GDPR)** and **AI-specific frameworks** are part of maintaining compliance and risk oversight in many organisations. Additionally, **regional policies** and **industry-specific standards** – such as **AI audits** and **responsibility standards** – should be incorporated into governance standards.

### Required security and privacy standards or certifications for AI technologies.

*% among total respondents, sorted descending.*



Once an organisation has outlined their responsible AI expectations and governance frameworks, next is establishing selection criteria for responsibly built AI technologies. These criteria should integrate the existing standards and focus on unique elements tied to generative AI such as transparency of origin, accuracy of outputs, training data licensing, bias mitigation and cultural localisation.

According to research findings, the top criteria organisations use when assessing generative AI technology include:

1. <b>Training data evaluation (72%)</b>	4. <b>Transparency of origin (55%)</b>
2. <b>AI use disclosures (63%)</b>	5. <b>Bias mitigation (50%)</b>
3. <b>Harm mitigation (60%)</b>	

These factors ensure that selected AI technologies meet both **business needs** and **ethical responsibilities**, supporting long-term organisational success.

Organisations should develop targeted selection criteria that align AI solutions with both strategic business objectives and responsible AI principles. These criteria emphasise:

<b>Transparency</b> Ensuring that AI processes are explainable and traceable.	<b>Cultural localisation</b> Adapting AI systems to respect diverse cultural and regional contexts.
<b>Accuracy</b> Maintaining high standards for data fidelity and predictive reliability.	<b>Bias mitigation</b> Actively reducing biases to support fair and equitable AI outcomes.

Documenting each stage of the assessment and selection process reinforces adaptability and accountability, creating a flexible governance model that can evolve with AI advancements and regulatory shifts. Below is a summary of steps to take in the assess phase:

## Assess

### Step 1: Evaluate organisational readiness.

- Define and communicate the company's standards on the responsible use of technology, inclusive of AI.
- CIO and/or cross-company committee review current systems and business processes to identify areas that will benefit most from responsible AI adoption.
- Aggregate input from internal business and functional leaders on additional use cases to consider for responsible AI adoption.

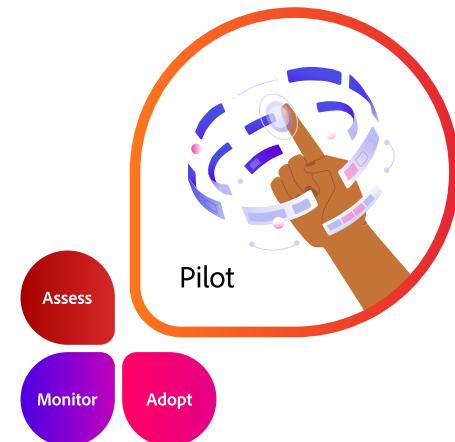
### Step 2: Select AI tech that is built responsibly.

- Review existing governance standards across privacy, security, accessibility and legal for AI considerations.
- Develop selection criteria that integrate the previously established standards that meet responsible AI expectations by focusing on transparency, accuracy, bias, cultural localisation and compliance.
- Evaluate and select AI technologies that best meet the established criteria and business needs, documenting the decision-making process.

## 2. Pilot: Identifying and piloting high-impact use cases.

The **pilot phase** bridges AI experimentation with operational reality. This phase allows key stakeholders to evaluate the technology's performance, in terms of how it aligns with business objectives and responsible AI goals. It goes beyond testing for technological feasibility, focusing on enabling key leaders and stakeholders to engage directly with the technology in meaningful ways. It is about enabling people to work with AI, make informed decisions on its scalability and ensure that it meets ethical, operational and regulatory standards.

Piloting gives organisations the chance to stress-test AI systems in context, helping them understand where accountability assessments and transparency documentation may be needed, as well as the performance of new capabilities relative to expectations. By documenting insights and gathering actionable learnings, organisations establish a roadmap for scaling AI responsibly, creating a foundation that supports both immediate and long-term objectives.



### 2.1 Identify and prepare priority use cases.

Developing a compelling AI business case includes engaging key stakeholders, front-line employees, to provide a holistic view of AI's potential. By involving those who will directly interact with the technology early on, organisations can identify high-impact use cases where AI delivers tangible benefits, such as in marketing content creation, coding, workflow automation and data management.

**Be specific – Focus on processes, not roles:** Rather than framing use cases around specific roles (e.g. "AI for developers"), focus on processes that AI can streamline and improve such as "AI-assisted coding for automating routine code reviews and error detection".

**Establish measurable metrics for usage and cost savings:** While ROI is important, AI pilots should also emphasise broader returns such as productivity, speed to market, employee satisfaction and enhanced customer experiences – metrics often referred to as "Return on Experience".

**Elevate impact beyond immediate gains:** Position the AI initiative as a driver of long-term transformation. Use cases should not only address immediate operational needs but also align with strategic goals like digital transformation or competitive differentiation.

### 2.2 Pilot against business and responsible AI criteria.

Evaluating pilots through a dual lens – business performance and responsible AI criteria – ensures that AI initiatives meet both operational goals and responsible AI benchmarks. Over half (54%) of the organisations surveyed have established an acceptable risk level for their priority use cases. Organisations should document these evaluations systematically, capturing learnings to inform future AI projects. This structured approach builds a robust foundation for scalable AI implementations.

### Actions when piloting:

**Set business and responsible AI benchmarks:** Define both operational goals (e.g. productivity, cost savings) and responsible AI metrics (e.g. transparency, fairness).

**Establish risk thresholds:** Set risk parameters and create a framework for ongoing assessments to manage and mitigate AI-related risks effectively.

**Capture and share learnings:** Develop a standardised process for documenting pilot outcomes to support transparency and guide future scaling efforts.

## Pilot

### Step 1: Identify priority use cases.

- From existing business priority use cases, identify 2–3 pilots where AI ethics and responsibility are important.
- For these use cases, establish metrics and thresholds to track both business and responsible AI performance.

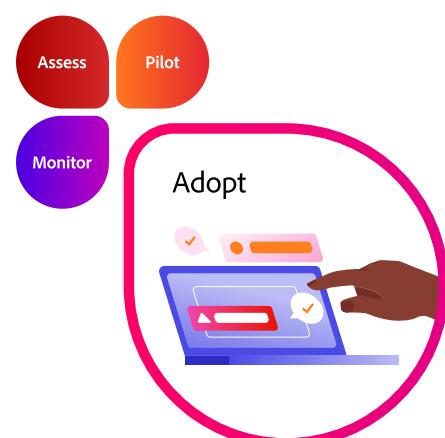
### Step 2: Pilot against business and responsible AI criteria.

- Execute pilots, with additional technical, business and responsibility validation and testing as needed.
- Evaluate pilot outcomes against pre-defined metrics and thresholds for business and responsible AI expectations, documenting learnings into future assessment and testing approaches.
- Advance to procurement/adoption based on pilot outcomes and insights.

## 3. Adopt: Integrating AI responsibly across the organisation.

The **adopt** phase marks the transition from pilot to organisation-wide integration. This phase focuses on transitioning from experimental applications to fully operational AI systems – deploying AI responsibly while embedding the lessons learned from pilots into real-world practices.

In this stage, your employees take active ownership of AI's role within their existing workflows. By building on their practical experience from the pilot phase, employees are equipped to drive AI adoption.



### 3.1 Train and enable the organisation.

Scaling AI effectively requires a knowledgeable workforce that understands both the capabilities of AI and the ethical responsibilities that come with its use. Tailored training programmes should help employees across roles and departments leverage AI tools. Many organisations (89%) recognise the importance of training, with nearly two-thirds including responsible AI guidelines. Training should integrate technical capabilities with principles of accountability, transparency and regulatory compliance.

#### Actions when training and enabling:

**Align training with governance:** Incorporate responsible AI guidelines into training materials to ensure that employees are aware of compliance, risk management and transparency requirements.

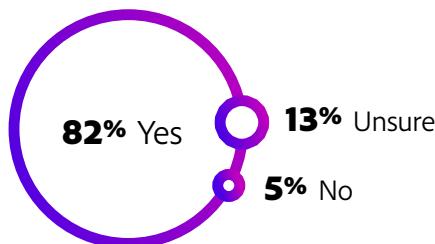
**Customise training to roles:** Develop tailored training modules addressing the needs of specific functions, including best practices for both business and responsible AI.

### 3.2 Deploy with responsibility in mind.

Adopting AI at scale requires establishing a governance framework that ensures responsible use. Organisations should align their AI initiatives with their existing governance policies, while continuously refining policies to meet evolving regulatory, operational and responsible AI standards.

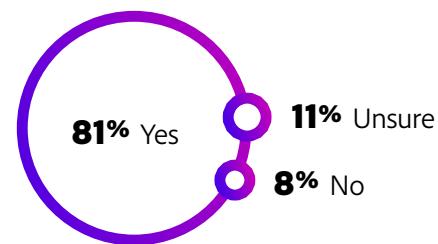
#### Plans to incorporate responsible AI considerations in standards for widespread deployment.

*% among total respondents, sorted descending.*



#### Inclusion of responsible AI considerations in technology governance efforts.

*% among total respondents, sorted descending.*



#### Actions for deploying responsibly:

**Foster a culture of accountability:** Encourage teams to understand the broader impact of AI on both operational workflows and stakeholder trust, instilling a sense of responsibility at every level.

**Continuously evolve training:** As AI governance frameworks evolve, update training programmes to reflect new best practices and regulatory changes.

## Adopt

### Step 1: Train and enable the organisations.

- Develop use case specific guidelines for how and when to use AI with responsible AI principles.
- Deploy comprehensive training to relevant groups as solutions are rolled out, inclusive of best practices.
- Celebrate and share wins across the organisation.

### Step 2: Deploy with responsibility in mind.

- For each technology, codify core adoption requirements (e.g. business impact, ease of integration and risk mitigation).
- Work closely with business leaders to align on trade-offs, where needed.
- Embed key AI and responsibility considerations into existing governance frameworks (e.g. access, control, roles).

## 4. Monitor: Continuous oversight and improvement.

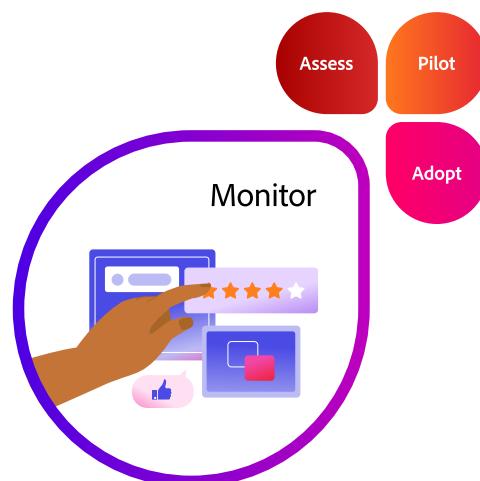
As AI systems move into full-scale deployment, continuous monitoring and improvement become essential.

The **monitor phase** emphasises real-time tracking, rigorous performance reviews and a proactive risk management approach to ensure AI systems remain effective, compliant and aligned with organisational goals. By embedding responsible AI metrics and establishing a structured review process, organisations can adapt to evolving regulatory demands and emerging risks while fostering sustained trust and operational value.

### 4.1 Monitor performance against business and responsible AI benchmarks.

Best-in-class monitoring of technology outcomes combines automated performance tracking with human expertise. While many organisations have adopted real-time monitoring tools to assess AI system performance, the effectiveness of these tools improves when human oversight is included. Human teams are better equipped to analyse data, spot risks and make informed decisions about necessary adjustments.

While 69% of organisations use real-time monitoring tools, these are significantly more effective when combined with human judgment. Many organisations prioritise technical metrics, with 72% focusing on accuracy and 69% on ROI, yet responsible scaling also requires attention to ethical dimensions. Embedding human oversight ensures transparency and predictability, building trust internally



and externally. Monitoring further enables proactive bias detection, with 49% of organisations tracking this metric, and 33% monitoring for harmful outputs. Without consistent, proactive monitoring, AI systems could compromise both integrity and trust. By refining performance monitoring to address both technical and ethical risks, companies protect their brand, build user confidence and lay a resilient foundation for scaling AI responsibly.

This collaboration between technology and people allows organisations to identify potential risks early, such as data inaccuracies, emerging biases or compliance lapses.

### AI-specific considerations for tracking technology performance and effectiveness.

% among total respondents

Accuracy of outputs **72%**

Performance of AI outputs (e.g. time savings, ROI) **69%**

Potential bias of outputs **49%**

Adoption rate of AI capabilities **47%**

Instances of harmful outputs **33%**

Other **2%**

None **5%**

### Nearly 1/3 of organisations

are not staffed to support continuous improvement on technology performance metrics and business outcomes

#### Other Mentions

- Being developed
- Data privacy related to model training
- BOT containment rate
- Abuse of insights

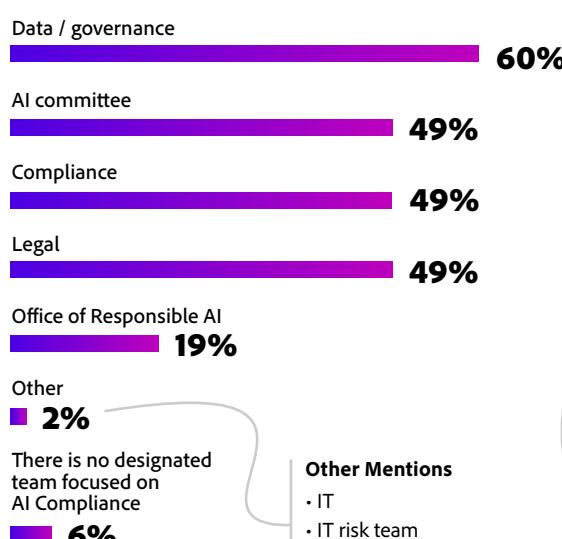
## 4.2 Ongoing risk management.

Risk management in AI is a continuous process that should evolve alongside AI systems. Establishing a structured, cross-functional approach to AI risk management enables organisations to proactively address both business and reputational risks. Scheduled reviews should involve stakeholders from across the organisation, including data scientists, business leaders and legal/compliance officers, ensuring a thorough evaluation of both technical performance and responsible AI goals.

AI risk management is a continuous process that evolves with technological advancements. Sixty percent of organisations involve data and governance teams, and 49% include AI committees, compliance and legal teams – highlighting the need for cross-functional collaboration. This proactive approach enables organisations to stay aligned with both internal values and external expectations. The “why” is about building resilience that adapts to regulatory shifts. With 68% of organisations emphasising responsible AI in risk management, comprehensive documentation and ongoing risk assessments are essential.

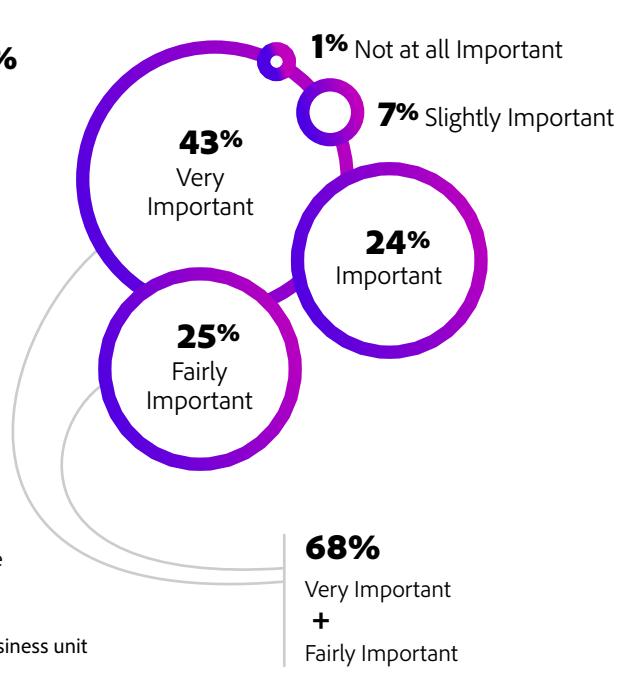
### AI-specific considerations for tracking technology performance and effectiveness.

% among total respondents



### Importance of responsible and ethical AI use considerations for teams ensuring compliance.

% among those who indicate a designated team focused on monitoring AI regulations and standards



By establishing rigorous performance metrics and fostering a culture of continuous risk management, organisations can ensure that AI deployments remain aligned with both business goals and responsible AI priorities. Proactive monitoring, combined with cross-functional reviews and thorough documentation, positions organisations to lead AI-driven transformation with both confidence and accountability.

## Key actions in risk management:

**Establish cross-functional risk reviews:** Create regular risk assessments involving data scientists, compliance officers and legal experts to identify emerging risks based on real-time data.

**Track and report consistently:** Gather continuous feedback from employees and end-users to detect usability issues, biases or unexpected behaviors.

---

## Monitor

### Step 1: Monitor performance.

- Define and track longer-term AI performance metrics, inclusive of business and responsibility goals.
- Establish ongoing review and discuss findings to continuously improve business performance while safeguarding the organisation's responsibility principles.

### Step 2: Deploy with responsibility in mind.

- Designate and empower roles to track evolving AI regulations and standards (e.g. Singapore AI-Verify, US Congressional proposals, etc.) and ensure company standards are updated accordingly.
- Develop process for continuous identification and mitigation of risk associated with the use of AI.
- Update documentation as to how organisation is adhering to company standards.

---

## ADOBE CASE STUDY

# Internal use of generative AI.

At Adobe, we see generative AI as transformational technology with the power to enhance human creativity, not replace it. We encourage responsible exploration of generative AI technology internally, aligned with our own AI Ethics principles of accountability, responsibility and transparency.

In June 2023, Adobe established a cross-functional internal working group sponsored by the CIO and CHRO, to help employees navigate the exploration and use of generative AI within Adobe in a safe, responsible and agile way. This group, working with leaders and subject matter experts across the company, focuses on guiding a thoughtful approach to grassroots employee experimentation by understanding the landscape of hypotheses for generative AI use, establishing appropriate guidelines and streamlining experimentation. The initiative has formed four persona-based working groups representing generative AI use cases across Adobe and has established an intake process, a generative AI risk tolerance framework and a blueprint for use case review which takes into account evolving ethical, security, privacy and other legal considerations. A list of approved generative AI tools and models based on specific use cases, and guidelines for employee use of generative AI are also available. The roll out of vendor generative AI guidelines in March 2024 included training sessions on using generative AI and the features in the chosen products.

Implementation of this initiative has helped streamline the process for faster experimentation and scaled application where possible, and enabled assessment of the companywide generative AI landscape. Adobe continues to foster shared learnings and insights across the business to create a collaborative ecosystem for collective exploration. The programme continues to evolve with the expansion of generative AI in our own products, proliferation of generative AI technology and models, as well as evolving legal and regulatory guidance. Experimentation review is monitored – the team is developing post-approval experiment tracking including for those experiments that go into production for scale.

## III. Embedding best practices in your organisation.

To adopt, monitor and optimise AI systems responsibly, organisations should focus on several operational areas: providing comprehensive employee guidance, rigorously evaluating vendors and establishing robust AI governance levers. By doing so, companies can ensure that their AI initiatives not only meet evolving regulatory standards but also build on existing governance and risk-management work while aligning with practices that foster trust, transparency and accountability.

This section outlines practical steps for embedding these best practices into your day-to-day operations.

### 1. Employee use guidance:

Tailoring AI usage guidelines to fit the specific needs and risks of your organisation is essential for ensuring responsible deployment. These guidelines should help employees navigate regulatory standards and governance protocols, aligning AI technologies with commitments to data security, transparency and accountability.

#### 1.1 Data sensitivity:

Clearly specify when data processing should occur locally or under strict access controls to prevent unauthorised access; this also means refraining from using prompts that could generate or manipulate sensitive outputs. This guideline protects proprietary information and maintains compliance with data privacy regulations.

#### 1.2 Transparency in AI usage:

Readily disclose the involvement of AI such as when it is used to create internal documents, customer facing interfaces or external communications. This practice fosters accountability and maintains trust in the authenticity and reliability of AI-generated content, maintaining a company's brand and reputation.

#### 1.3 Account management policies:

Establish explicit policies for the use of generative AI tools that require account registration, including defining whether organisational email accounts can be used, specifying which tools are approved for business purposes and discouraging use of personal accounts for work content. This safeguards against unauthorised use and helps maintain alignment with the organisation's broader information security programme.

By tailoring these guidelines to your organisational context, employees can navigate the use of generative AI tools confidently and responsibly, contributing to an environment where innovation and integrity go hand in hand.

## 2. Vendor evaluation: Sample questions.

Evaluating AI vendors requires informative questions and understanding what answers you are looking for to ensure that their systems adhere to responsible AI, legal and regulatory standards. The following questions are designed to provide a baseline assessment, enabling organisations to make informed decisions about potential partnerships and mitigate risks associated with AI adoption.

Theme	Vendor question	Rationale	Adobe example
<b>Data provenance and usage</b>	<i>"What specific types of data were used in the development and training of the AI system?"</i>	Provides insight into the sources, nature and scope of the data used to train AI models. This question ensures that the vendor's data practices align with buyers' responsible AI standards and comply with legal objectives.	<b>Firefly:</b> Adobe does not include enterprise user content (including Firefly inputs and outputs) in datasets used to train Firefly foundation models.
<b>Intellectual property compliance</b>	<i>"Were any datasets utilised that may have copyright, intellectual property or licensing restrictions?"</i>	This inquiry verifies that all data sources are legal and obtained through approved mechanisms, preventing potential legal disputes.	<b>Firefly:</b> Customers can review licence history information at any time at <a href="https://stock.adobe.com/Dashboard/LicenseHistory">stock.adobe.com/Dashboard/LicenseHistory</a> while logged in with their Adobe credentials.
<b>Training data and logic transparency</b>	<i>"Can you provide a detailed explanation of the training data and the logic applied in developing the AI system?"</i>	Transparency in these aspects identifies potential biases and provides understanding of the model's reasoning process, which is critical for evaluating its reliability and fairness.	<b>AEP AI Assistant:</b> Adobe does not use any customer data to train or fine-tune the Azure OpenAI service.
<b>Output clarity</b>	<i>"Can you provide a plain language description of the AI system's outputs?"</i>	Ensuring that outputs are understandable to non-technical reviewers allows for effective decision-making and responsible usage.	<b>Firefly:</b> Adobe automatically generates <a href="#">Content Credentials</a> for certain Firefly-generated assets to help provide transparency that the asset was created using Generative AI.
<b>Human oversight</b>	<i>"If human review is part of the AI system, what is the extent and nature of human involvement?"</i>	Knowing the balance between automated processes and human judgment evaluates the system's operational dynamics and identifies areas where human intervention may be necessary to uphold quality and responsibility standards.	<b>Acrobat AI Assistant:</b> Adobe strictly limits who can access this information to a small number of trained Adobe employees directly involved in the development of the Adobe Generative AI Service.
<b>Fairness and bias evaluation</b>	<i>"How was the AI system assessed for bias and what were the outcomes of these evaluations?"</i>	This question demonstrates vendor's commitment to equitable AI practices and their methods for detecting and mitigating biases that could affect different demographic groups disproportionately.	<b>Acrobat AI Assistant:</b> Adobe teams conduct testing to reduce the potential for biased and harmful outcomes in our generative AI products. See <a href="#">Generative AI Built for Business solution brief</a> .
<b>Risk mitigation</b>	<i>"Has there been an assessment of potential harmful outputs, and what measures were implemented to mitigate these risks?"</i>	Assessing a vendor's proactive steps in identifying and addressing potential negative outcomes shows harm remediation. This means that the AI system is tested for operating safely and responsibly.	<b>AEP AI Assistant:</b> Adobe uses internally developed content filters to (a) determine if the input (prompt) in AI Assistant in AEP adheres to Adobe's Generative AI User Guidelines and (b) filter out any generated responses that violate these guidelines (e.g. hate speech and profanity).

By posing these targeted questions, organisations can evaluate AI vendors and make informed choices that align with their AI responsibility standards and operational priorities. This approach manages risks and ensures that AI vendor relationships are built on a foundation of transparency, compliance and responsible innovation.

### 3. AI governance levers.

Implementing robust AI governance ensures that AI systems are developed, deployed and monitored in a way that aligns with organisational values and regulatory standards. There are many regulatory standards in place such as the European Union AI Act and frameworks such as Singapore's AI Verify. In the U.S., companies should follow state-comprehensive privacy laws and the NIST AI Risk Management Framework as it is the probable foundation of future regulation.

The following governance levers can help organisations manage AI risks and enhance transparency, accountability and security:

<b>AI inventories</b>	Establish inventories of AI systems, categorising them according to risk profiles and strategic priorities to serve as a centralised repository.	Have you documented your AI use case and categorised relevant risks?
<b>Feedback mechanisms</b>	Establish robust feedback channels for capturing insights from those outside the AI development team, such as end-users, customers or the public.	What feedback channels are being used to capture insights from end-users, customers or the public?
<b>System limitations documentation</b>	Document AI system limitations including information on the AI model's knowledge gaps and the context in which its outputs can be reliably used.	Have you documented known or expected limitations for your AI use cases?
<b>Content provenance</b>	Trace and verify the origin, history and modifications of AI-related data, including tracking training data sources, algorithms used and transformations.	How are you tracking origin, history and modifications of AI-related data, including data sources and transformations from creation to end use?
<b>AI testing and red teaming</b>	Assess risks such as the unintentional exposure of training data, susceptibility to reverse engineering and risks associated with model extraction.	What testing protocols have been pursued and how do they address specific AI use case risks?
<b>Secure software development</b>	AI systems should be integrated into the organisation's secure software development lifecycle, adhering to established best practices for coding and deployment.	How has your AI use case integrated existing secure software development protocols?
<b>Training</b>	Training should cover relevant policies, procedures and compliance requirements, equipping stakeholders with the knowledge to manage AI risks effectively and act in accordance with organisational standards.	Have you participated in AI governance and risk management training?

Embedding each of these best practices into existing processes ensures the responsible adoption, management and optimisation of AI technologies. By integrating these activities such as providing robust employee guidance, evaluating vendors rigorously and establishing comprehensive AI governance levers, organisations can create a resilient baseline of practices that aligns with governance standards and regulatory requirements, equipping them to not only meet today's challenges but also to scale AI responsibly for future success.

## IV. Responsible implementation builds responsible innovation.

Maximising the potential of AI in a company requires sourcing responsibly built technology, establishing clear usage guidelines, developing dedicated training and deploying strong governance. This approach drives business value and ensures that AI initiatives meet regulatory expectations and uphold responsible implementation standards, embedding a culture of responsible AI throughout the organisation.

Consistent oversight and adaptation keep AI initiatives on track. By defining performance metrics, conducting regular evaluations and proactively managing risks, organisations can stay ahead of regulatory shifts and maintain the integrity of their AI projects.

Moving into the future with AI, this framework equips organisations to take a leading role within the responsible AI landscape. With a focus on impact, integration and integrity, this approach paves the way for sustainable innovation and enduring success.



© 2024 Adobe. All rights reserved.

Adobe, the Adobe logo and Acrobat are either registered trademarks or trademarks of Adobe in the United States and/or other countries.