onetrust | AI Responsible Artificial Intelligence Institute

Securing reliable AI solutions: Strategies for trustworthy procurement

EBOOK | 2024

onetrust

AI Responsible Artificial Intelligence Institute

AUTHOR CREDITS:

Lauren Diethelm, AI Governance Content, OneTrust.

Bex Evans, Senior Product Marketing Manager, OneTrust.

Hadassah Drukarch, Director of Policy & Delivery, Responsible AI Institute.

Patrick McAndrew, Member Engagement & Community Manager, Responsible AI Institute.

DISCLAIMER:

No part of this document may be reproduced in any form without the written permission of the copyright owner.

The contents of this document are subject to revision without notice due to continued progress in methodology, design, and manufacturing. OneTrust LLC shall have no liability for any error or damage of any kind resulting from the use of this document.

OneTrust products, content and materials are for informational purposes only and not for the purpose of providing legal advice. You should contact your attorney to obtain advice with respect to any particular issue. OneTrust materials do not guarantee compliance with applicable laws and regulations.

Table of Contents

Building a foundation of trust in AI procurement

With 97% of organizations actively engaging with AI and 74% incorporating Generative AI (GenAI) technologies in production, AI is as much a business imperative as it is a technological breakthrough. Trust in these systems is crucial to their successful integration and operation. Yet, this trust is often compromised by a lack of shared understanding between organizations and their third-party vendors regarding responsible AI practices.

Many organizations rely on third-party vendors to integrate AI systems into their operations, procuring entire AI systems from designers and suppliers or selecting certain design elements to add to their own in-house systems.

However, these partnerships can fall short in establishing a shared understanding of responsible AI, which can compromise AI governance and negatively impact both business operations and the users who place their trust in AI services.

As most companies depend on external vendors for their AI solutions, procurement becomes the key to unlocking the full potential of responsible AI. By proactively addressing the challenges associated with AI procured through third parties, organizations can build the necessary foundation to transition toward AI systems that prioritize trust, transparency, and accountability. This guidebook offers a roadmap and key considerations for organizations looking to procure trustworthy AI systems.

The what and the why behind trustworthy AI

Since the launch of ChatGPT in November 2022, AI adoption has surged dramatically, with ChatGPT reaching 100 million users in just two months, setting a record for the fastest-growing consumer application. This rapid growth has driven AI advancements globally, leading to new models, widespread adoption, regulatory progress, and cutting-edge research. However, real-world applications reveal gaps that can lead to unpredictable consequences beyond human oversight.

Trust is crucial for AI to reach its potential. Yet, a 2023 global study by KPMG found that 61% of people across 17 countries hesitate to trust AI.

Trust is the currency of innovation, stemming from regulation, certifications, and standards.

As AI continues to evolve, ensuring its reliability and safety will be key to gaining public confidence and maximizing its benefits.

While eliminating all risks associated with AI adoption is impossible, organizations must implement strategies to mitigate short, medium, and long-term risks. With AI governance frameworks consistently lagging behind technological developments, organizations face challenges in developing guardrails and ethical policies for AI's development and use.

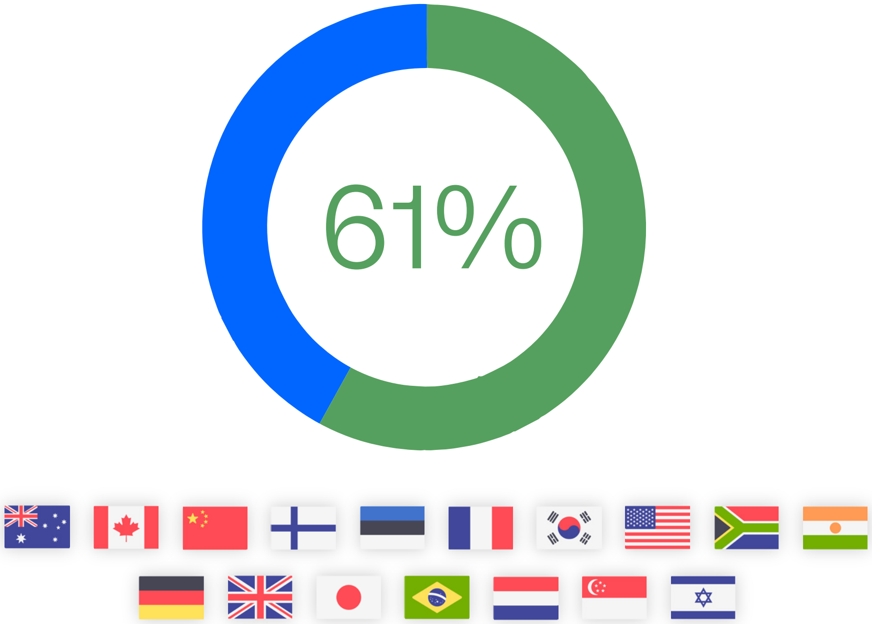

Despite these challenges, recent AI advancements have spurred regulatory initiatives from governments, NGOs, and the private sector. This includes the EU AI Act, the 2023 US Executive Order on Trustworthy AI, NIST's AI Risk Management Framework (AI RMF), and the ISO/IEC 42001:2023 standard for Artificial Intelligence Management Systems. These regulations focus on characteristics such as safety, security, explainability, privacy, fairness, reliability, and transparency – essential for measuring AI trustworthiness.

Assessing these characteristics requires human judgment and balancing based on the AI system's context of use. Regulations and standards are vital for setting healthy boundaries for AI, ensuring responsible development, use, and procurement, thereby fostering trust and reliability.

A roadmap for procuring trustworthy AI from third parties

I. AI strategy – business problem, risk appetite, controls

As organizations observe the widespread adoption of AI among their peers and consumers, they may feel an urgency to adopt AI themselves. However, even though AI development is accelerating daily, that doesn't necessarily mean it's the appropriate solution for every problem.

Before determining whether your organization is procuring AI from a trustworthy vendor, it's worth evaluating if an AI solution is the right choice at all or whether a different solution would better enable you to perform the same tasks and achieve the same objectives – without the inherent risks that AI brings.

As a starting point, you must ask if there's a legitimate AI use case and, if so, clarify whether the potential value generated is worth the risk and cost of investment.

To answer these key questions, begin by defining the legitimate business problem you're looking to solve with AI and document a clear use case for the AI solution you are looking to procure.

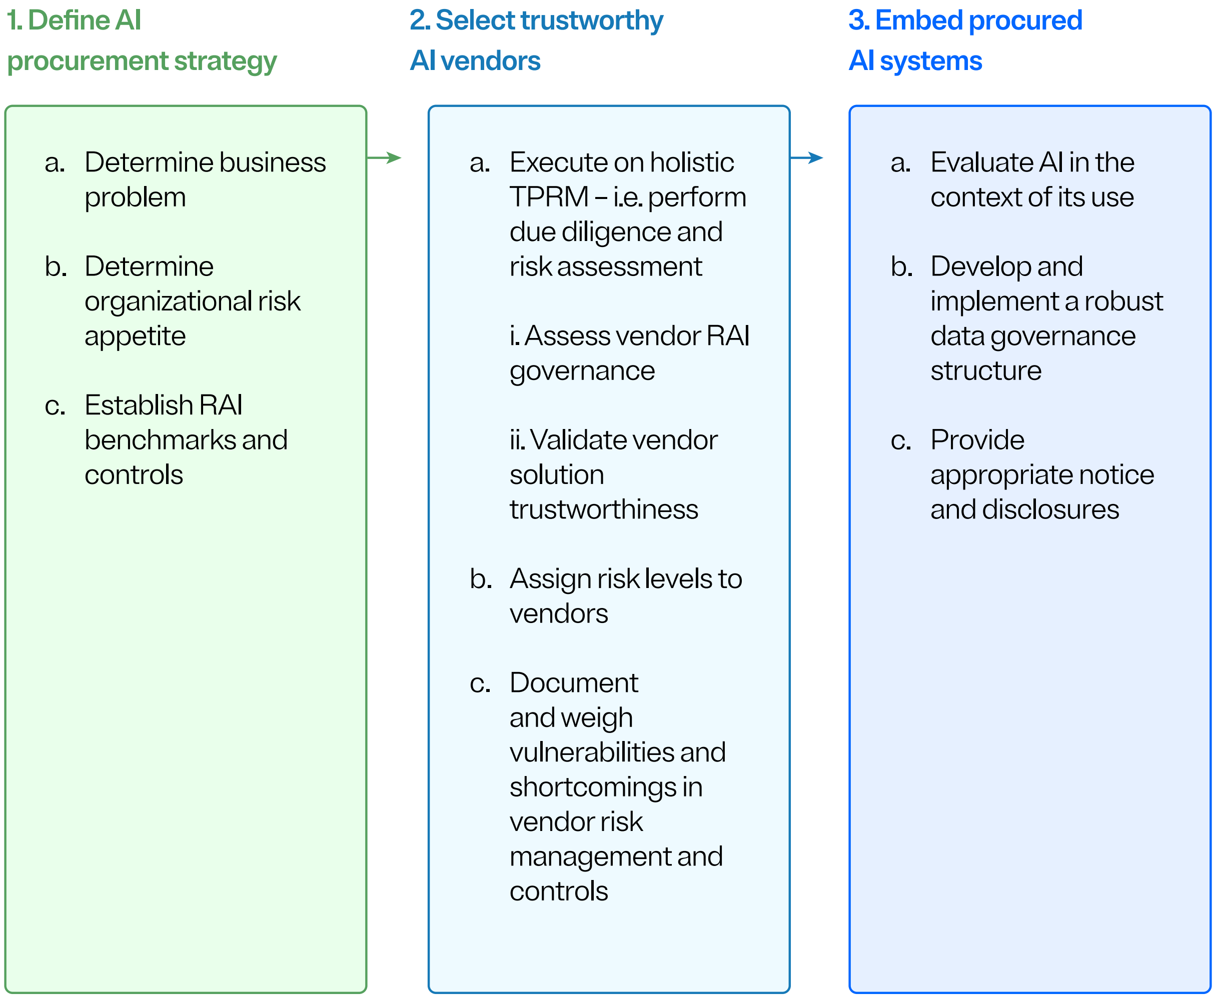

1. Define AI procurement strategy

a. Determine business problem

b. Determine organizational risk appetite

c. Establish RAI benchmarks and controls

2. Select trustworthy AI vendors

a. Execute on holistic TPRM – i.e. perform due diligence and risk assessment

i. Assess vendor RAI governance

ii. Validate vendor solution trustworthiness

b. Assign risk levels to vendors

c. Document and weigh vulnerabilities and shortcomings in vendor risk management and controls

3. Embed procured AI systems

a. Evaluate AI in the context of its use

b. Develop and implement a robust data governance structure

c. Provide appropriate notice and disclosures

A roadmap for procuring trustworthy AI from third parties

This exercise will help you determine whether AI is the right solution for your specified business problem or whether a cheaper (and less risky) alternative is available.

McKinsey has identified three common archetypes for companies using AI: takers, or users of available tools; shapers, who integrate available models with proprietary data; and makers, or builders of large language models.

Depending on which archetype best describes your AI usage, the total cost of that project will vary. McKinsey estimates that for takers simply procuring and using an available model, the cost can range from $0.5 million to $2 million. Makers developing their models might see their investment reach as high as $200 million – not counting ongoing maintenance costs.

When defining the business problem you hope to solve with AI, consider the additional value you expect to generate. For instance, if you're looking to implement a customer success chatbot for potentially upward of $2 million, is that chatbot poised to generate that much value for your business?

Outside of the startup cost, the increased risk accompanying AI will increase your cost as well. Different risk controls need to be implemented, costing resources and working hours as you monitor and maintain your system.

Whether or not this risk is worth the cost will depend entirely on your organization. Factors like your industry, your organization's risk appetite, the use case, and the type of data used by the AI system will all affect whether the risk of the AI system is worth the cost of implementing it.

This is not to suggest you shouldn't use AI in your business. Rather, it's to illustrate that there's increased risk (and increased cost) when creating and using large-scale, custom AI models. The more personalized your AI becomes – as you move further from being a taker toward being a maker – the more expensive it becomes to build and maintain.

This direct relationship between customization and cost will drive organizations toward a hybrid approach that leverages embedded technologies or third parties.

Given this trend toward using third-party vendors (or a more shaper approach), having a robust process for vetting and implementing trustworthy vendors is crucial.

A roadmap for procuring trustworthy AI from third parties

II. Selecting trustworthy AI vendors

Organizations looking to responsibly harness the potential of AI within their operational workflow should select an AI vendor that balances innovation with appropriate risk management.

Broadly speaking, organizations face several challenges when procuring AI solutions. Firstly, developing organizational responsible AI capacity is an essential step. This involves training and upskilling teams on AI ethics, governance, and best practices to ensure they are equipped to evaluate and integrate AI responsibly.

Secondly, navigating legal uncertainty is a significant hurdle due to the rapidly evolving nature of AI technology and regulatory incentives, which often outpace the development, harmonization, and implementation of regulatory frameworks. Organizations must stay informed on regulatory changes and ensure their procurement strategies are adaptable, thereby determining and adequately prioritizing compliance practices.

Finally, addressing transparency challenges in AI procurement is crucial. Organizations must demand clear documentation and explanations from vendors about how their AI systems work, including data sources, algorithms, and decision-making processes, to ensure alignment with responsible AI standards and organizational values.

Organizations should adopt a holistic approach to third-party risk management (TPRM) and clearly communicate it to potential vendors. This approach must evolve with the integration of AI and align with responsible AI standards. It should build upon existing strategies to prevent governance silos and ensure that benefits to individuals and society are at the core of the procurement process.

Procurement teams serve as gatekeepers, responsible for shaping responsible AI adoption through rigorous due diligence and risk assessment. It’s their responsibility to bind vendors to transparency and accountability while engaging in cross-functional collaboration with diverse stakeholders.

Performing a responsible supplier assessment is a key step in this process, breaking down the due diligence and risk assessment process into assessing a vendor’s AI governance processes and practices, on the one hand, and validating a vendor’s AI solution trustworthiness, on the other.

In line with this risk-based approach, organizations should first assess vendors’ AI governance practices across different dimensions of organizational maturity in alignment with responsible AI standards – such as ISO/IEC 42001:2023 and the NIST AI RMF – focusing on strategy and leadership, policy and governance, people, training, and resources, tools and processes, and procurement practices.

At this stage, the procurement team might request evidence of AI governance processes, defined roles and responsibilities for AI development, ethical AI conduct codes, and adherence to AI governance standards.

A roadmap for procuring trustworthy AI from third parties

Secondly, vendor solutions should be scored across their lifecycle stages and according to the system's risk level: low, medium, and high. Such an assessment may differ depending on the context and the chosen procurement approach (be it licensing, integration, or building the solution) and should include, at minimum, considerations related to accountability and transparency, safety, privacy, fairness, explainability and interpretability, reliability, and robustness.

The outcomes of this holistic TPRM approach form the basis for informed vendor selection, aligning with the organization's risk appetite and proactively identifying vulnerabilities and shortcomings in vendor risk management and controls.

By embedding these steps into their procurement strategies, organizations can confidently navigate the complexities of AI adoption, balancing innovation with risk management to achieve their strategic objectives and build and maintain trust in the procurement ecosystem.

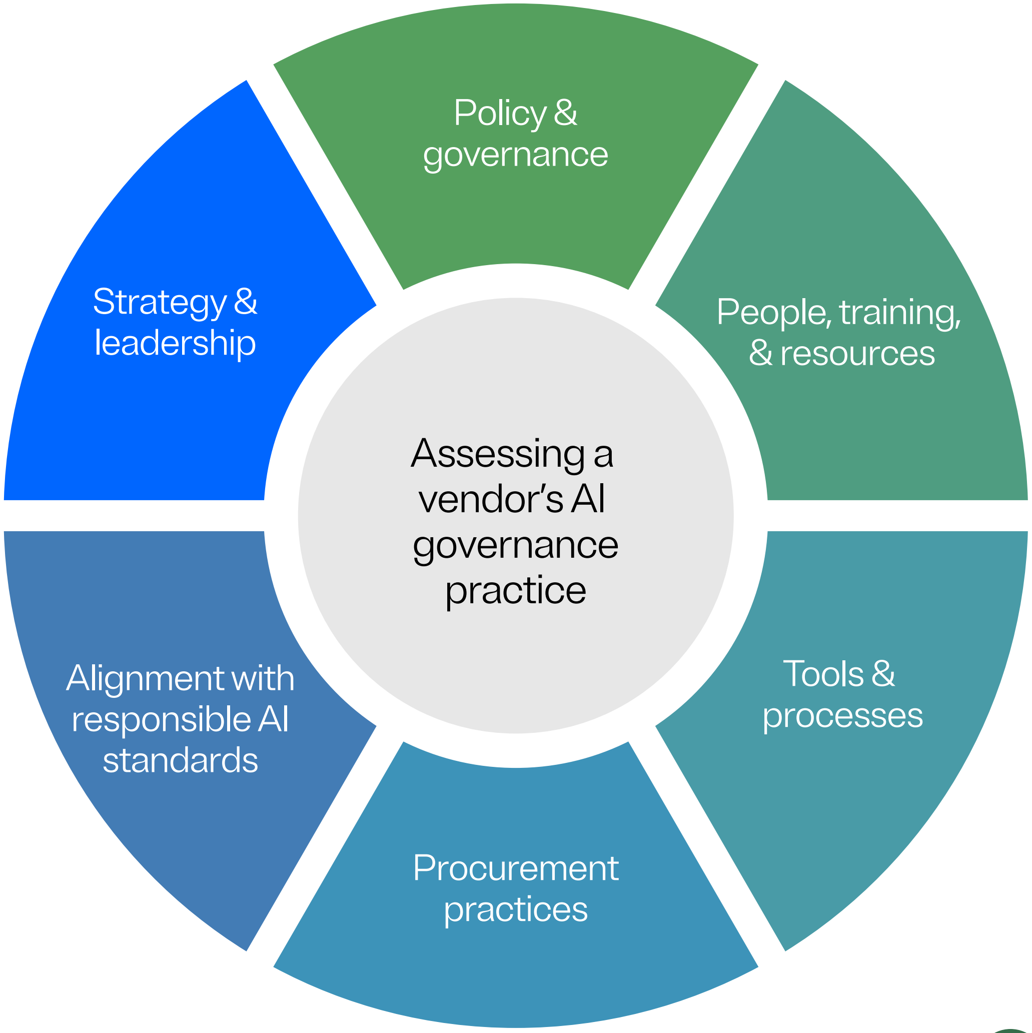
Policy & governance
People, training, & resources
Tools & processes
Procurement practices
Alignment with responsible AI standards
Strategy & leadership
Assessing a vendor's AI governance practice

A roadmap for procuring trustworthy AI from third parties

III. Embedding AI systems

Though many organizations currently fall somewhere in McKinsey's shaper archetype, preparing for third-party AI risks could better equip businesses to handle in-house AI development risks in the future. But whether you're embedding a procured AI system or developing and deploying your own, there are a few key considerations to keep in mind.

First, context is king. Regardless of whether an AI system is from a third party or developed in-house, it's important to evaluate the technology within the context of its use; both the EU AI Act obligations for deployers and implications of recent FTC rulings support this idea.

There will always be implications based on the data used by an AI system. Regardless of whether you're embedding an existing system, procuring and adapting one, or developing your own, each of those routes has different implications for the data being used. Depending on your role in each of those scenarios, you may have additional requirements for protecting that data.

Although your obligations to protect individuals' data and privacy don't disappear with AI, managing third-party AI risks may differ from developing AI in-house. This requires specialized skills in data science, machine learning, and model development. What new program components must be considered as an organization transitions from using an AI system to deploying it?

Having a robust data governance structure in place will help ensure you're protecting the data you're using, regardless of where the AI model originated. Training data is always necessary for AI, whether it's for an externally available model or an in-house ML model.

Protecting this data is vital to ensure responsible AI practices, and embedding AI-specific controls early in the software development lifecycle can help you do this.

Lastly, disclosures and transparency are equally important whether you're procuring third-party AI or developing your own model. It's essential that end users are clearly informed when they are interacting with an AI system. This transparency not only builds trust but also aligns with emerging regulatory requirements that emphasize the need for clear communication about AI's role in decision-making processes.

Turning strategy into action for trustworthy AI adoption

Organizations are racing to embrace AI, specifically GenAI, with Gartner predicting that more than 80% of enterprises will have used or deployed this technology by 2026, up from just 5% in 2023. Despite the high expectations companies have for the value that GenAI can add to their businesses, Bain & Company has noted that a clear vision of how to create practical business value from this technology is still largely lacking.

A recent study from McKinsey, which analyzed 1,000 companies, reveals that companies with leading digital and AI capabilities outperform slower companies by two to six times in shareholder returns across every sector analyzed. Moreover, the gap between leaders and laggards is widening as the effective implementation of digital and AI brings compounding benefits. Simply put, the opportunity cost of AI non-adoption is accelerating.

Before procuring AI, organizations must carefully evaluate if AI is the right solution for their business problem. This involves defining clear use cases, assessing potential value against risks and costs, and understanding the organization's risk appetite.

A holistic third-party risk management approach is crucial when selecting AI vendors. Whether procuring third-party AI or developing in-house solutions, organizations must evaluate AI systems within their context of use, implement robust data governance structures, and ensure transparency and clear disclosures to end users about AI interactions.

If your organization is lagging behind competitors in AI adoption, prioritizing AI use will buy time to build the collaborative muscle necessary to tackle game-changing innovations and identify program redundancies.

By following this guidance and implementing actionable and transparent measures, organizations can better navigate the complexities of AI adoption while balancing innovation with responsible risk management.

Turning strategy into action for trustworthy AI adoption

To successfully procure and implement AI systems, organizations must:

- Conduct thorough due diligence on vendors' adherence to responsible AI best practices.
- Implement a governance plan to manage the risks associated with third-party AI systems.
- Create a vendor assessment framework incorporating responsible AI principles.
- Implement ongoing monitoring and evaluation processes for AI systems.
- Invest in training and upskilling teams on responsible AI and governance.
- Stay informed about evolving AI regulations and standards.
- Be proactive in addressing potential vulnerabilities or shortcomings in vendors' risk controls by establishing appropriate risk mitigation measures for AI system deployment and operation.
- Prioritize transparency and accountability in all AI-related activities.

By embedding these controls into your procurement strategy, your organization can confidently embrace AI, ensuring it contributes to your strategic objectives while maintaining trust and accountability in the procurement ecosystem.

About

onetrust

About OneTrust

OneTrust unlocks the full potential of data and AI, responsibly. Our platform enforces the secure handling of company data, empowering organizations to drive innovation responsibly while mitigating risks. With a comprehensive suite of solutions spanning data and AI security, privacy, governance, risk, ethics, and compliance, OneTrust enables seamless collaboration between data teams and risk teams to enable rapid and trusted innovation. Recognized as the market leader in trust, OneTrust boasts over 300 patents and serves more than 14,000 customers globally, ranging from industry giants to small businesses.

Onboard AI vendors confidently – get started with OneTrust today.

AI Responsible Artificial Intelligence Institute

About Responsible AI Institute

Founded in 2016, Responsible AI Institute (RAI Institute) is a global and member-driven non-profit dedicated to enabling successful responsible AI efforts in organizations. We accelerate and simplify responsible AI adoption by providing our members with AI conformity assessments, benchmarks and certifications that are closely aligned with global standards and emerging regulations.

Navigate AI procurement challenges confidently and turn them into opportunities for innovation. Become a Responsible AI Institute member to access expert guidance, tools, and frameworks that will elevate your organization's approach to responsible AI adoption.

onetrust | AI Responsible Artificial Intelligence Institute