Deloitte.

Operationalizing AI governance

MAKING AN IMPACT THAT MATTERS since 1845

Contents

Introduction

Today's organizations must navigate an ever-evolving, increasingly complex ecosystem of AI applications that drive critical business outcomes and span the entire AI lifecycle. Even the AI sector itself seems poised at the edge of disruption, as recent model developments call into question existing business and technical assumptions. However, no matter the AI tools and models used, organizations must tackle the same critical challenge: how to scale AI responsibly and in a way that builds trust, ensures regulatory compliance, and delivers operational agility.

To achieve sustainable, responsible growth amid fragmented data landscapes, siloed processes, and emerging regulations such as the EU AI Act, organizations need to establish comprehensive AI governance frameworks and policies—and operationalize them using next-gen technical platforms. This means seamlessly embedding AI governance principles and oversight across every stage, from the initial design and development of AI systems to their deployment, ongoing monitoring, and optimization. Doing so enables organizations to ensure the AI systems they use are transparent, accountable, and aligned with ethical principles and their own values.

This guide describes how decision-makers can equip their organizations with the tools needed to scale AI with confidence, foster stakeholder trust, and drive sustainable innovation in an increasingly regulated environment. We provide an overview of the critical elements required to operationalize AI governance and identify the key challenges organizations face in doing so. We explore practical capabilities and governance models, and offer real-world examples of implementing AI governance. Finally, we offer recommendations on how organizations can begin their own journey towards building scalable and trustworthy AI systems.

Insights from the industry landscape

Growing diverse vendor capabilities
There is no one vendor that addresses all aspects of AI governance and trustworthy AI. Organizations need to adopt a partnership-based approach (with ecosystem partners and vendors) in operationalizing AI governance frameworks.

Enablement and maturity towards change management
Organizations are gearing towards enablement and focusing upskilling and change management, by evolving the capabilities and how they will be deployed across the enterprise.

Ongoing effort on integration of processes and tools
There has been a requirement to standardize the development process by seamlessly integrating the next-gen technical platforms and AI Governance processes, enabling coherence through a robust structure.

Enhanced role of AI centre of excellence
Institutions are leveraging the AI centre of excellence (CoE) function to act as a catalyst of innovation, to foster partnerships (e.g., academia, vendors) and lead with trust.

Evolving AI/GenAI regulatory landscape
AI regulations are here, and institutions need to ensure that guardrails are being placed as they are investing and scaling AI.

Key challenges in operationalizing AI governance

Establishing effective AI governance is a significant undertaking. There are several key challenges that must be overcome as organizations implement governance frameworks that ensure AI development fosters trust, enables agility, and aligns with global regulatory standards.

Data landscape and data quality issues

The data landscape most organizations must deal with is complex and fragmented, marked by data silos, burdened by legacy systems, and suffering from poor data quality. This landscape hinders AI adoption and exposes organizations to data-related risks regulatory scrutiny. Strengthening data governance is critical to mitigating such risks and facilitating responsible AI deployment.

Underdeveloped, rigid governance frameworks

Many organizations' existing governance frameworks are too underdeveloped, outdated, and inflexible to easily adapt to the demands of the rapidly evolving AI landscape. The lack of resources for AI governance councils and other critical capabilities makes it hard for organizations to develop strong AI governance and respond effectively to new regulations.

Lack of alignment between AI projects, business goals, and risk processes.

Organizations' AI projects often lack clear goals and objectives, and it can be difficult to see how these projects align with broader business strategies. This presents a significant challenge for AI governance, making it harder to deploy resources efficiently and measure project success. It also prevents organizations from fully realizing the benefits of their AI projects—and hinders their ability to manage AI-related risks effectively.

Lack of effective guardrails throughout the AI lifecycle

AI tools can produce biased or otherwise unreliable outputs unless the AI lifecycle includes guardrails designed to protect against this. Unfortunately, organizations may lack these guardrails, hindering AI governance efforts.

> "Organizations are struggling to land and articulate business cases with clear and tangible ROI of investment in GenAI.
>
> **Jas Jaaj, Global GenAI Ecosystems & Alliances Leader**
> **Deloitte Canada**

Choosing the right vendor for your organization

Vendors serve as the "connective tissue" for AI governance platforms, and organizations must build a collaborative ecosystem and form strategic partnerships with the right vendors to scale AI responsibly. Yet given the sheer number of vendors in the market, it can be difficult for organizations to select the vendor(s) with whom they can work seamlessly to ensure interoperability, effective governance, and responsible AI deployment. The complexity of integrating disparate vendor solutions and internal systems is often exacerbated by organizations' lack of native tooling.

Organizations can address these challenges through an AI governance platform—a comprehensive solution for managing data landscapes, ensuring regulatory compliance, articulating business cases, and testing guardrails. Integrated seamlessly into existing and planned processes, AI governance platforms can also help address vendor and tooling challenges best, ensuring compatibility across diverse tools and managing vendor compliance with governance policies. Leveraging AI governance platforms enables organizations to enhance AI governance, mitigate AI risks, and accelerate AI adoption and impacts.

Making AI governance real: practical capabilities and governance models

Organizations need three practical capabilities, or foundational elements, to operationalize AI governance: a clear strategy driven by AI leaders, robust governance framework(s), and the appropriate technology to manage the program efficiently and scale it globally. In the illustration below, a hypothetical organization's AI strategy and governance is guided and executed though an AI centre of excellence that brings together executive sponsors, decision-makers, advisors, and execution-focused leaders and teams. The governance process is based on Deloitte's Trustworthy AI framework, complemented by relevant regulations in other jurisdictions. An AI governance platform provides the technology needed to monitor and manage AI systems and workflows.

1 AI centre of excellence

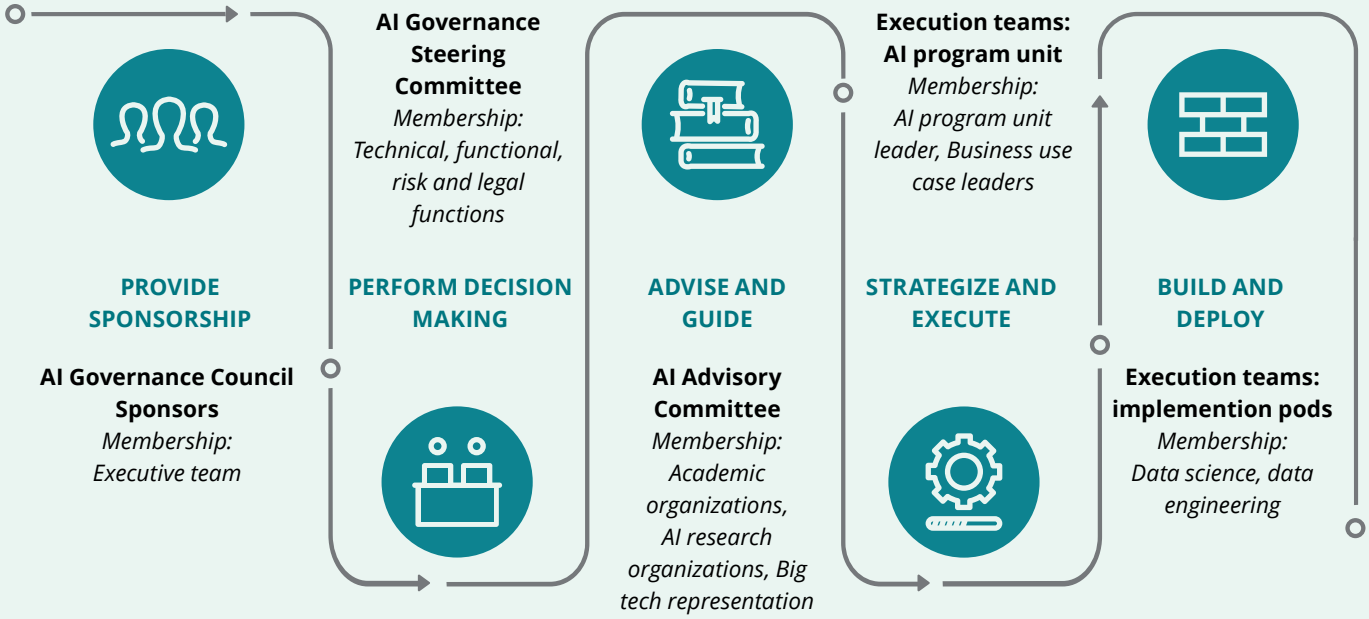

2 Enterprise AI risk framework

Complemented by regulatory frameworks and compliance requirements (e.g., EU AI Act, NIST)

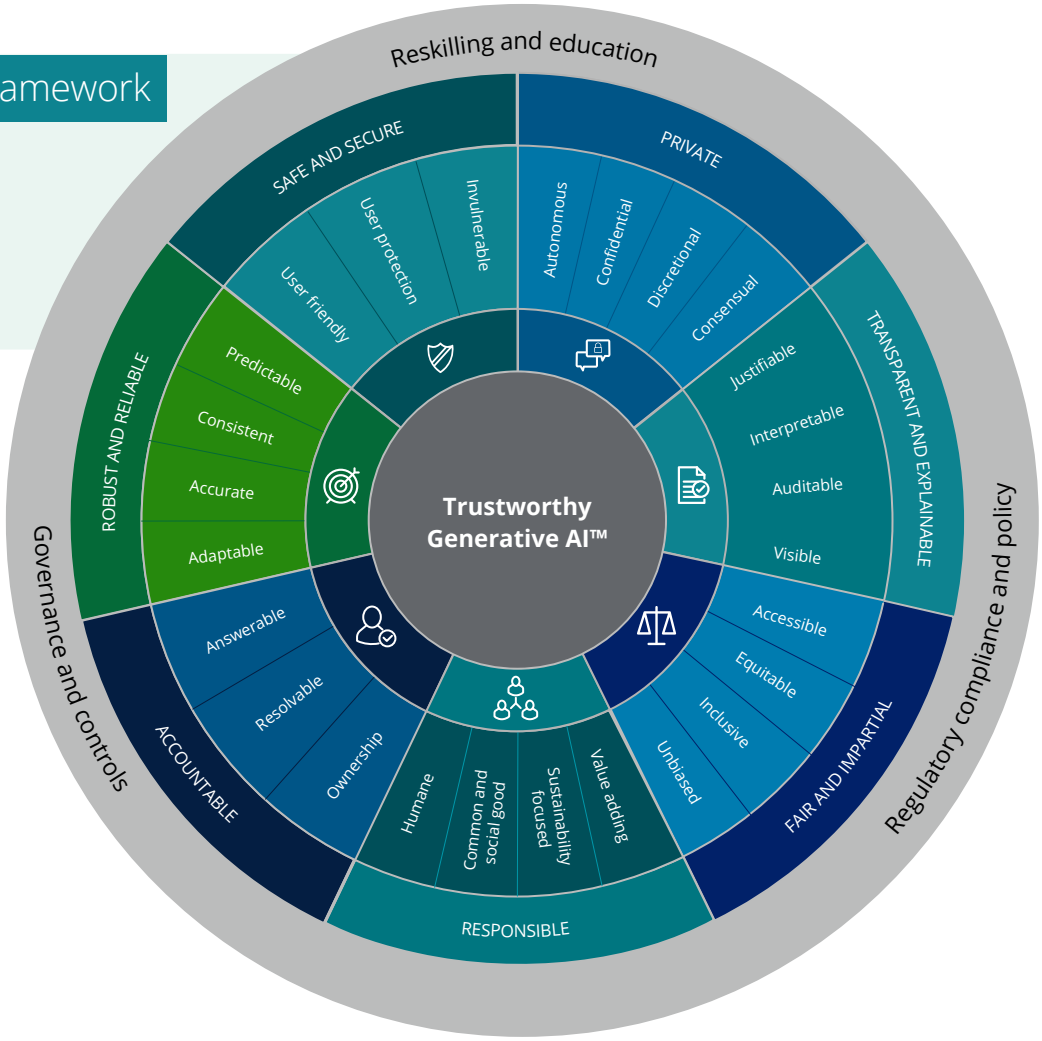

3 AI governance platform

Technology enablement to:

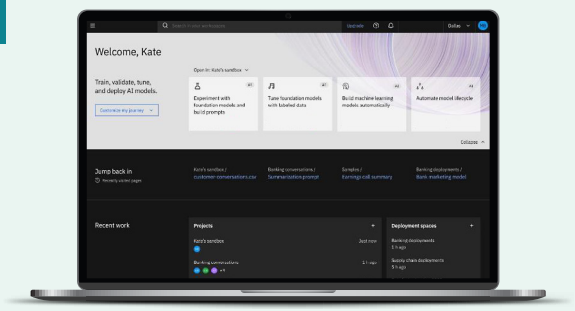

View and manage the AI inventory
All use case, models, and datasets

Monitor AI systems
Flag breaches in near real-time and respond (e.g., drift, bias, toxicity, etc.)

Evaluate AI risk, manage workflows
Across multiple lines of defense

Provide transparency and earn
Support reporting and compliance

> Having proper governance solutions in place allows organizations to elevate confidence, move at speed, and reuse a lot of existing solutions that have been deployed across different geographies and departments.

Jas Jaaj, Global GenAI Ecosystems & Alliances Leader,
Deloitte Canada

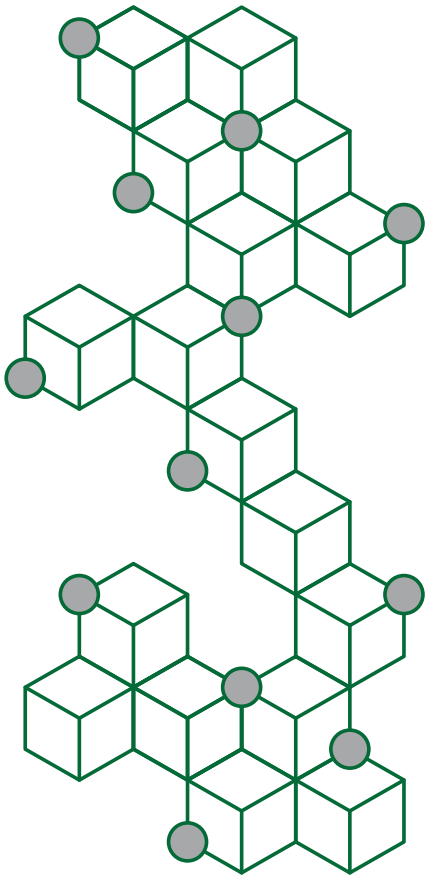

A hybrid AI governance model delivers both consistency and flexibility

A hybrid AI governance model that combines centralized oversight with local flexibility enables an organization to achieve global consistency, while preserving its ability to adapt to diverse regulatory and operational needs.

In this hybrid approach, a centralized in-house clearing hub validates AI models to ensure they comply with organizational principles, ethical standards, and regulatory requirements before deployment. Governance, risk, and compliance (GRC) tools streamline workflows, approval processes, and IT service management (ITSM) integration, ensuring consistency across AI operations. Model risk management (MRM) establishes model inventories, performs risk assessments, and ensures AI reliability to foster trust and mitigate enterprise-level risks.

At the same time, the hybrid AI governance model delivers the flexibility needed to adapt centralized standards to local contexts. AI observability and traceability enables real-time monitoring of AI systems to ensure transparency, accountability, and alignment with local regulations, while Ecosystem-based frameworks equip regional teams to implement tailored intake systems, monitoring, and reporting to meet specific, local operational requirements while staying aligned with centralized standards.

Ultimately, adopting a hybrid approach to AI governance empowers organizations to scale AI responsibly while remaining agile in a dynamic regulatory landscape.

Implementing AI governance: real-world examples

Driving GenAI innovation in financial services through AI governance

A century-old financial services organization and private provider of financial retirement services was committed to embracing AI and driving GenAI innovation responsibly. The organization engaged Deloitte to support them on their journey to drive the development and implementation of GenAI use cases in line with its "AI-first" mindset.

To facilitate this, Deloitte supported the establishment of an AI Centre of Excellence that would catalyze GenAI innovation throughout the enterprise. Deloitte conducted an independent review of the organization's policies and governance frameworks, providing GenAI-specific enhancements, considerations, and suggestions. Additionally, Deloitte supported the organization in redesigning their model development life cycle (MDLC) process to accelerate GenAI use cases' time to market and time to insights.

By implementing these guardrails, the organization was able to safely drive GenAI innovation through a collaborative approach that aligns key stakeholders across risk domains. This ensures that all GenAI use cases are developed and implemented in a trustworthy manner.

Operationalizing AI governance at Deloitte

Deloitte has long leveraged our own Trustworthy AI framework to develop platforms and accelerators that leverage AI, machine learning, and GenAI. These platforms and accelerators, leveraged in client engagements across our global member firms, increasingly require the use of AI and GenAI technologies to provide improved capabilities that drive differentiated outcomes.

However, AI and GenAI capabilities also bring the potential for new risks and necessary compliance with emerging regulations and legislation. As a result, Deloitte's global member firm leadership directed platform ownership teams to select a "system of record" to deliver operationalized AI governance across critical client platforms and accelerators. Critically, the selected platform needed to not only provide comprehensive GRC and MRM capabilities, but also provide:

- End-to-end monitoring for AI use case development and deployments
- A "single pane of glass" dashboard functionality to align different member firms and jurisdictions
- Flexible deployment architecture to support an array of platform infrastructure and deployment typologies.

Ultimately, Deloitte selected watsonx.governance—a full AI lifecycle toolkit developed by IBM—as the "system of record" because it best provided AI governance capabilities (e.g., risk detectors, usage monitoring, model management) within a scalable deployment architecture (e.g., OpenShift deployment architecture, ITSM integration, and deep hyperscaler integration).

Key takeaways when operationalizing AI governance

1 **There is no one-size-fits-all solution:** Not all tools provide a complete AI governance solution. Organizations should focus on interoperable, transparent, and adaptable platforms that align with business needs and enable scalability and compliance.

2 **Alignment is key to scaling AI responsibly:** Structured partnerships, clear roles, and a strong operating model are vital for scaling AI responsibly and with speed. Effective governance aligns AI initiatives with organizational goals, ensuring trust and compliance.

3 **Standardize AI, ML, and GenAI processes:** Embedding governance into the AI/ML lifecycle ensures consistency and accountability. Clear workflows and standardized processes help organizations address operational and regulatory requirements efficiently.

4 **Governance platforms should provide an integrated view:** It's important that governance platforms embedded in data pipelines offer an end-to-end operational view. This enables seamless data integration, real-time insights, and effective regulatory alignment.

5 **Continuous monitoring is essential:** Built-in observability and transparency metrics allow organizations to detect anomalies, refine models, and maintain accountability. Continuous monitoring ensures AI systems remain trustworthy and can adapt to evolving regulations.

6 **Governance platforms/tools must be as flexible as the deployment architecture:** Organizations' deployment environments and legacy IT systems influence the scope and breadth of AI governance solutions. If an organization's underlying deployment environment is multifaceted and complex, it's critical that any AI governance platforms or tools must provide capabilities to support that environment. Having a flexible AI governance model is also essential to enable organizations to adapt to future developments in a fast-changing AI landscape.

How to get started with operationalizing AI governance

As AI models, tools, and solutions evolve at incredible speed, the demand for AI governance has never been greater. More and more organizations are turning to AI to achieve their operational and strategic objectives—and as they do, it's critical that they proactively establish the building blocks needed to operationalize AI governance.

Ready to move forward with AI governance but not sure how to begin? Here are some key considerations:

- Define and identify the core components of your AI governance council: sponsors, cross-functional steering committees, and execution teams.
- Understand your current capabilities related AI governance and, as required, prioritize upskilling and training to help ensure readiness for AI governance execution.
- Define best practices and establish trustworthiness before deploying AI and GenAI further.
- Identify AI governance use cases aligned to your organization's business problems.
- Evaluate AI governance platforms based on their technical functionalities and deployment capabilities as well as how well they address your identified challenges and adhere to industry compliance standards.
- Establish an AI governance ecosystem that aligns platform, technology, and operating model(s) to help provide full lifecycle understanding/monitoring and identify opportunities to better realize value from your AI, ML, and GenAI use cases.
- Define AI governance workflows, identify accountabilities for AI governance operationalization, and put in place necessary controls for compliance.
- Deploy comprehensive model risk management to ensure data transparency in model development.
- Monitor AI-related regulatory developments regularly and assess potential required improvements to your model or governance strategy.

Contact us

Ready to begin your AI journey? Deloitte can help. Reach out today to start the conversation:

Preeti Shivpuri
Partner, Trustworthy AI Leader
pshivpuri@deloitte.ca

Jas Jaaj
Global GenAI Ecosystems & Alliances Leader
jjaaj@deloitte.ca

John Mantia
Global IBM Alliance Manager, Global GenAI Leader
jmantia@deloitte.com

We wish to thank the following Deloitte people for their contributions to this report.

Atif Jaleel
Manager, Office of Generative AI

Sarah Lee
Senior Consultant, Data Risk

Ingrid Leung
Consultant, Data Risk

Harvey Singh
Senior Consultant, Data & AI

Aditya Sharma
Consultant, Financial Crimes Practice

Deloitte.

www.deloitte.ca