Bridging policy and practice: AI governance for Indian enterprises

Data governance knowledge series – Topic 7

January 2026

Contents

01

Why AI governance matters for your organisation's future

As AI continues to evolve and broaden its usage across sectors, organisations are increasing their investment in AI technologies across their business functions. Although their adoption is becoming more commonplace, PwC's 2026 AI Business Predictions highlights a difference in business outcomes as a result of how organisations embrace AI. While most companies see modest AI gains, a select few achieve transformative outcomes such as surging top-line growth, valuation premiums, and operating or business model reinvention.[1]

These successes often share two common characteristics—disciplined execution and strategic focus. Disciplined execution calls for a high level of governance across various steps in AI implementation. Organisations need to focus on developing the right AI governance framework that will facilitate as well as protect the data assets of the organisation as per internal policies and regulatory practices around AI. When disciplined execution is coupled with strategic implementation in key business lines, it can propel the organisation into a transformatory trajectory powered by AI.

AI governance is effective when it is purposeful and comprehensive. It is expected to ensure compliance with evolving regulations, mitigate any risk to the enterprise, and build trust. In order to turn AI from a tactical tool into one of competitive advantage, AI governance must address three drivers: regulatory acceleration, enterprise risk exposure, and competitive advantage and trust.

Figure 1: Drivers of AI governance

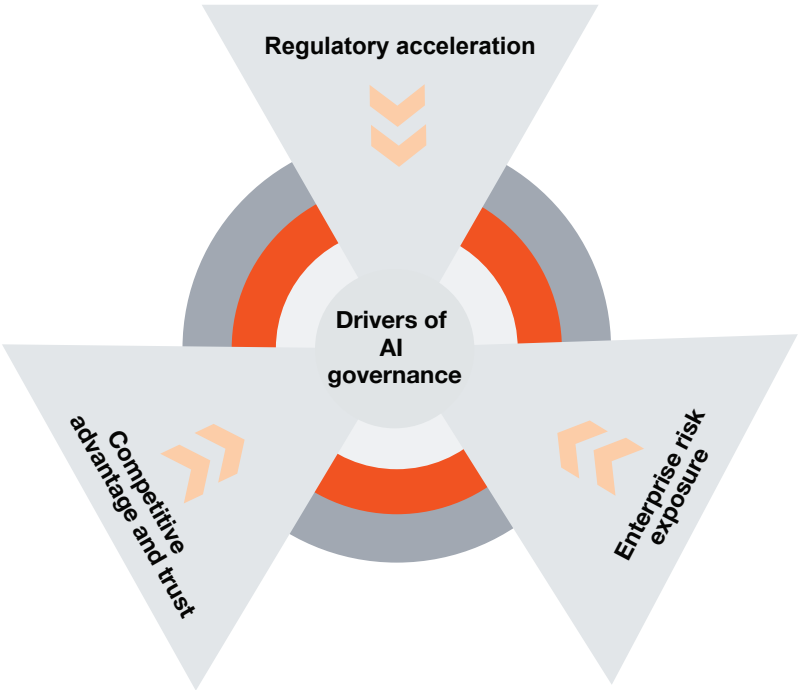

1 https://www.pwc.com/us/en/tech-effect/ai-analytics/ai-predictions.html

1. **Regulatory acceleration:** Countries around the world are crafting AI governance frameworks converging on risk-based, principle-driven governance with stronger controls for high-impact uses and AI models to strike a balance between protection of rights and pro-innovation industrial strategy.

Figure 2: Key frameworks shaping AI governance across the seven regions

EU's Artificial Intelligence Act

Risk-based framework with stringent requirements for high-risk AI, aiming for transparency and accountability

Canada's Artificial Intelligence and Data Act (AIDA)

Focus RAI use, consumer protection, and fair competition

Brazil's Bill No. 2,338/2023

Risk categorisation of AI systems with compliance requirements to ensure secure, reliable AI

Israel's Artificial Intelligence Regulations and Ethics

Principle-based, sector-specific approach using non-binding ethical guidelines and voluntary standards to encourage responsible innovation

Japan's promotion of AI-related technologies

Innovation-focused framework with non-binding guidelines, establishment of an AI Strategy Centre, and emphasis on voluntary compliance and international cooperation

New Zealand's Algorithm Charter for Aotearoa New Zealand

Commitments for fair, ethical, and transparent public-sector algorithm use, emphasising human oversight and Māori data sovereignty

South Africa's National AI Policy Framework

Human-centred AI strategy addressing socio-economic disparities through talent development, digital infrastructure, and ethical governance

2. **Enterprise risk exposure:** As per our playbook on 'Managing the risks of generative AI', the technology holds potential to expose an organisation to four broad risks—data risks, model and bias risks, prompt or input risks, and user risks.[2]

Figure 3: The risks of generative AI explained

Data risks	Model and bias risks	Prompt or input risks	User risks
Error propagation, intellectual property (IP), or contractual issues (due to lack of approvals to use data for such purposes), or misleading and harmful content caused by low-quality data used to train generative AI models	Breach of ethical and RAI principles in the language model development, leading to discriminatory or unfair outputs	Misleading, inaccurate, or harmful responses due to unsophisticated prompts or questions being provided to the AI model	Unintended consequences due to users becoming unwitting parties to the creation of misinformation and other harmful content. For instance, they might pass off AI-generated hallucinations—erroneous or nonsensical responses—as fact.

Source: Managing the risks of generative AI

Having structured governance embedded into an organisation's AI lifecycle—from ideation to deployment and monitoring—helps mitigate such risks. Without structured governance, AI models may leak sensitive company information and expose the organisation to significant compliance risks. AI models may also introduce bias and discrimination, model drift, and privacy violations causing business failure and loss of customer trust.

3. **Competitive advantage and trust:** With governance, organisations can ship AI models faster with fewer mistakes, with features that organisations, regulators, and customers value—such as safety controls to prevent harmful outputs, explanations of how the models worked on their answers, and source citations. These features, in addition to being ethical, also act as trust drivers which is a key differentiator in today's competitive market.

Ultimately, this means that organisations may choose to pivot from seeing governance as a compliance routine or a chore, to viewing it as a strategic lever—using it to unlock the full value of AI for them and their customers.

In this context, India's AI Governance Guidelines propose a framework that strikes a balance between AI innovation and responsibility, as well as advancement and safety. It is important to note that this article does not delve into responsible AI (RAI), which focuses on embedding principles of fairness, transparency, and safety into AI models and products. However, through AI governance, enterprises can lay a foundation that enables them to operationalise the principles of RAI.

2 https://explore.pwc.com/generativeai?_pfses=pcuxtVAzgkCGqTrUvVuiWjNm

02

Understanding India's vision for AI governance

In November 2025, the Ministry of Electronics and Information Technology (MeitY), under the IndiaAI Mission, released the India AI Governance Guidelines.[3] These guidelines detail a foundational, principle-based governance framework designed to balance innovation with responsibility.

The central tenet of the guidelines are the **seven 'sutras', or guiding principles**, as highlighted in Figure 4.

Figure 4: The seven guiding principles of India's AI Governance Guidelines

Trust is the foundation

Innovation and adoption will stall in the absence of trust.

People-first

Human-centred design, human supervision, and human empowerment

Innovation over restraint

When all else is equal, responsible innovation ought to take precedence over cautious restraint.

Fairness and equity

Prevents discrimination to encourage inclusive development

Accountability

Enforcing rules and assigning responsibilities clearly

Understandable by design

Ensuring that users and regulators can comprehend the disclosures and explanations

Safety, resilience, and sustainability

Robust, safe systems that are resilient to systemic shocks and sustainable in terms of the environment

3 https://static.pib.gov.in/WriteReadData/specificdocs/documents/2025/nov/doc20251156856O1.pdf

In order to operationalise these principles, the guidelines enumerate six governance pillars, as shown in **Figure 5.**

Figure 5: Six pillars as per India's AI Governance Guidelines

01

Infrastructure

Support adoption by providing access to high quality data sets (for evaluation), reliable computing resources and integrating with digital public infrastructure.

02

Capacity building

Start training, education, and skill-building initiatives to empower individuals, foster trust, and raise knowledge of the benefits and risks associated with AI.

03

Policy and regulation

Adopt policies that are flexible, agile, and balanced, and examine existing legislation to foster innovation and reduce the dangers associated with AI.

04

Risk mitigation

Create a risk assessment system tailored to India that considers actual evidence of harm.

05

Accountability

Establish a tiered liability structure according to the function carried out, the degree of risk, and the observance of due diligence.

06

Institutions

Adopt a whole-of-government strategy in which public bodies, sectoral regulators, and ministries collaborate to create and execute AI governance frameworks.

These guidelines will have significant impact on how enterprises in India develop and deploy AI systems. Enterprises will have to pay attention to ensure their practices abide by the framework. This would in turn help them achieve the following:

1. Obtain trust from their customers for adoption of RAI practices.
2. Gain a competitive advantage over their competitors in the same sector.
3. Mitigate risks of any AI-related harm caused by AI models.

The next question for organisations is clear: How can we translate the national-level principles into internal enterprise-level governance? We address this in the next section.

03

Building an effective AI governance framework for Indian enterprises

To effectively utilise the guidelines within enterprises, it is important to address two key questions: what to govern and how to govern.

What to govern: Understanding the scope

In order to design a meaningful AI governance framework in alignment with the national guidelines, it is important to consider the scope of AI governance. This governance is not limited to the just AI models being deployed, but spans the entire AI lifecycle, including the data ecosystems that feed, host, and operationalise these AI models.

- AI lifecycle: AI models evolve through a series of phases—from strategy and planning to deployment and operation[4] to model retirement. Each phase carries risk, dependencies, and governance requirements and therefore, enterprises need to consider governing the entire lifecycle.
- Data ecosystems: AI is only as trustworthy as the data it is trained on, or that which it consumes. Therefore, it is important to ensure oversight of data pipelines, storage systems, and applications from which models get their input. Amongst databases, data lakes, and external data sources, it is important to govern modern AI architecture stores such as feature stores, vector databases, and prompt or response stores in the case of generative AI.

4 https://www.pwc.de/en/risk-regulatory/responsible-ai/ai-lifecycle-management-as-part-of-ai-governance.html

How to govern: Building the structures and mechanisms

After establishing the facets in an enterprise which is recommended to be governed, organisations can proceed to operationalise the six pillars of India's AI Governance Guidelines.

In this section, we present an indicative set of actions that organisations can take to structure their governance approach under each pillar.

Pillar	Indicative actions
Infrastructure This pillar extends to access control, data lifecycle management, and the readiness of environments in which AI training and deployment happen.	1. Create an inventory of the data sets, environment of operations, and AI models including their risk rating. Record and analyse the risk associated with the use of any third-party systems or data in the landscape. 2. Create data quality, data security standards and ensure transparency and explainability across the AI lifecycle through documentation.
Capacity building This pillar suggests enterprises should consider developing competence of not just technical capability but also awareness of ethical considerations, risks and obligations, and governance expectations.	1. Create training plans and conduct regular training for all roles associated with the AI lifecycle—data engineers, data scientists, data governance teams, system owners, and business departments. 2. Provide AI governance and training courses organised by sectoral bodies and government institutions to the workforce.
Policy and regulation This pillar suggests that AI models may present risks that overlap with regulatory, legal, or internal privacy and security considerations for enterprises.	1. Document and map regulatory, legal, and enterprise privacy and security guidelines against the functional and technical aspects of the AI model to identify any gaps. Discuss measures to tackle the identified gaps and associated risks. 2. Ensure that documentation of the AI model includes the principles of transparency, its adherence to legal obligations, and risk assessments.
Risk mitigation This pillar advises to develop risk assessment and classification frameworks as per organisational context. This supports in assessing and mitigating risks across the AI lifecycle.	1. Create risk assessment, classification frameworks, and incident reporting mechanisms to identify and classify AI risks and use cases and to track and analyse the risks. Create a contingency plan or grievance handling mechanism to handle any failures that might occur from an erroneous AI model. 2. Create testing approaches aligned with the risk assessment framework for testing the AI applications prior to deployment. 3. Create tests to determine if bias exists across all AI lifecycle stages that may cause harm to individuals, groups, or society.
Accountability This pillar emphasises on human-led accountability and clarification on how different stakeholder groups across the AI lifecycle are governed.	1. Create a RACI matrix for all stakeholders involved in the AI lifecycle and ensure that all roles and responsibilities are understood. 2. Design forums to coordinate and manage risks among stakeholder groups in the AI development lifecycle.
Institutions This pillar primarily addresses line ministries, sectoral regulators, standards bodies, and public institutions to lead in setting standards, ensuring collaboration and enforcing compliance.	1. Organisations must cooperate with sectoral agencies and regulators on specific policies or guidelines brought forward for compliance. 2. Organisations can consider participating in events conducted by the government and sectoral bodies to remain updated on any upcoming AI regulations.

04

Illustrating AI governance: Building a cohesive framework

Let's consider a use case for an Indian life insurance provider offering long-term savings and protection products. During the process of claim filing, customers must submit multiple documents including identity proofs, medical reports, hospital bills, death certificates, and police reports as applicable. The claims teams manually verify these documents to ensure alignment with policy terms, beneficiary identity, and rider conditions—making the process time-consuming and prone to errors. The objective of AI implementation here would be to automate this verification process for speed, accuracy, and transparency in claims handling.

The implementation approach would broadly involve the following:

1. An automated document processing pipeline to extract key data points from documents
2. A smart validation mechanism that can help to validate the claim and check consistency with the policies
3. A human-in-the-loop workflow integration to review and add any missing information

If an AI governance approach had to be built around this implementation, the following considerations would have to be made:

Figure 6: Considerations for effective AI governance

Analyse datasets and integrations

- Identify data categories. Establish guidelines for data minimisation, security, and retention as per DPDP Act 2023.
- Map data flows across integrated systems, assess control gaps, and implement role-based access controls for data access.

Assess risk levels

- Assess risk levels of document type and validation rule.
- Define confidence thresholds for automated decisions and create a human-in-the-loop workflow to ensure accuracy in decision making.

Use governance forums

- Discuss the AI model progress, data usage, and compliance with data owners, data stewards and system owners in the AI governance forums.
- Communicate updates or decisions made on the usage of data.

Enable explainability mechanisms

- Implement the AI model to provide natural language summaries of the decisions taken by the model.
- A confidence score can help claim handlers gauge the need for a manual review.

Implement data security controls

- Ensure that the data security principles of the organisation including access controls and data storage are taken into consideration for the AI model as well.

Validate and monitor

- Once implemented, conduct periodic reviews to verify how compliant the system is with organisational, legal, and regulatory policies.

The way forward for Indian enterprises

AI has now become a tangible and transformative force shaping markets, customer expectations, and regulatory landscapes. The imperative is clear: to avoid the risk of being outpaced, organisations that utilise AI with disciplined execution and strategic focus will achieve speed, innovation, and trust.

By aligning enterprise AI governance frameworks with India's AI Governance Guidelines, enterprises can translate national principles and pillars to actionable organisational structures across their lifecycle—from strategy, planning, and deployment to model operation and retirement.

The time to act is now. Organisations that proactively embed governance are not just avoiding penalties but are better positioned to create resilient, trustworthy, and differentiating AI products in today's competitive landscape.

About PwC

We help you build trust so you can boldly reinvent

At PwC, we help clients build trust and reinvent so they can turn complexity into competitive advantage. We're a tech-forward, people-empowered network with more than 364,000 people in 136 countries and 137 territories. Across audit and assurance, tax and legal, deals and consulting, we help clients build, accelerate, and sustain momentum. Find out more at www.pwc.com.

PwC refers to the PwC network and/or one or more of its member firms, each of which is a separate legal entity. Please see www.pwc.com/ structure for further details.

Contact us

Mukesh Deshpande
Partner, One Consulting
PwC India
mukesh.deshpande@pwc.com

Indranil Mitra
Partner, One Consulting
PwC India
indranil.mitra@pwc.com

Rajnil Mallik
Partner, One Consulting
PwC India
rajnil.mallik@pwc.com

Prakash Suman
Managing Director, One Consulting
PwC India
prakash.suman@pwc.com

Authors

Abhishek Chaurasia

Tanvi Yerrapragada

Reemal Prabhod

Editorial

Rashi Gupta

Design

Shipra Gupta

pwc.in

Data Classification: DC0 (Public)

SG/January 2026 - M&C 50515