BOARD BRIEF SERIES | AI GOVERNANCE

NEXORA TECH

The Twelve Questions Every Board Must Ask Before the Next AI Investment Goes Wrong

AI Governance is no longer a technology agenda item. It is a fiduciary responsibility — and the questions boards are not asking today will become the liabilities they defend tomorrow.

By Aparna Kumar
Founder — Nexora Tech | Former CIO — SBI & HSBC | Board Advisor

April 2026

Executive Summary

AI is moving from the IT agenda to the boardroom table — and not always at the board's choosing. When a regulator issues a notice, when an automated decision harms a customer, or when an investor asks why your AI disclosures cannot be substantiated, the question is no longer theoretical. It is existential.

Boards across India and globally are confronting a disquieting reality: the pace at which AI has been adopted within organisations has, in many cases, materially outpaced the governance structures built to oversee it. Artificial intelligence now touches credit decisions, KYC processes, fraud flagging, pricing models, customer service interfaces, and hiring pipelines. Yet in boardroom after boardroom, the AI agenda is still treated as a management briefing rather than a standing governance item.

This Board Brief is not written as a technology primer but as a governance document. It is designed to equip Independent Directors, Audit and Risk Committee members, and ESG Committee chairs with the conceptual architecture and practical questions they need to exercise meaningful oversight — not to manage AI, but to govern the enterprise that does.

Four Strategic Imperatives Demanding Board Attention Now

Regulatory Exposure

India's Digital Personal Data Protection Act, RBI's IT governance circulars, and SEBI's growing interest in algorithmic accountability are creating a compliance environment where AI-related failures carry direct board-level consequences.

AI Disclosure Risk

The SEC has already charged investment advisers for misleading statements about AI capabilities — so-called AI washing. Indian regulators are watching. Boards that cannot substantiate their public AI claims face legal, reputational, and valuation consequences.

Consumer Harm Liability

The Air Canada chatbot ruling established a landmark principle: organisations are accountable for AI outputs presented to customers, even when those outputs are wrong. Indian enterprises that deploy customer-facing GenAI face the same exposure.

Capital & Valuation at Stake

PE and institutional investors are beginning to embed AI governance due diligence into their investment frameworks. Enterprises that cannot demonstrate structured AI oversight will face valuation discounts and restricted access to growth capital.

The twelve questions structured throughout this Board Brief are not a compliance checklist. They are a governance architecture — informed by OECD accountability principles, NIST's AI Risk Management Framework, and the emerging ISO/IEC 42001 management system standard. They are the questions that distinguish a board that governs from one that merely receives reports.

Context & Urgency Framing

The Velocity Problem

Generative AI has compressed what would ordinarily take a decade into less than 3 years. Boards that approved AI pilots in 2022 are now presiding over production deployments embedded in core business processes — often without a corresponding upgrade to governance infrastructure. The vocabulary has changed faster than the oversight mechanisms.

> What distinguishes this inflexion point from previous technology waves is the nature of AI's decision-making footprint. Unlike ERP systems or digital banking platforms, AI systems make — or materially influence — decisions that directly affect customers, employees, and counterparties.

Credit refusals. Fraud flags. Insurance pricing. Content moderation. These are not operational efficiencies. They are consequential actions that carry legal, ethical, and reputational weight.

2022

Boards approve AI pilots

2024–25

Production deployments embedded in core business processes

Today

Governance infrastructure has not kept pace — oversight mechanisms lag the vocabulary

The Regulatory Environment

India's regulatory posture on AI and digital governance is tightening on multiple axes simultaneously:

DPDPA

The Digital Personal Data Protection Act, once fully operationalised, will impose obligations around automated decision-making that directly intersect with AI systems. Data fiduciary responsibilities will fall squarely within the board's governance remit.

RBI IT Governance

RBI's IT governance frameworks and outsourcing guidelines extend implicit expectations around algorithmic risk management to regulated entities deploying AI in financial processes.

SEBI

SEBI's growing interest in algorithmic accountability — particularly in trading and investment advisory contexts — signals an expanding perimeter of regulatory attention.

Globally, the trajectory is equally clear. The EU AI Act establishes a risk-based classification regime with significant obligations for 'high-risk' systems. The EU Digital Operational Resilience Act (DORA) embeds AI-related operational risk into financial sector governance requirements. The SEC's AI washing enforcement actions have established disclosure standards that are influencing regulatory thinking well beyond the United States.

Why Boards Cannot Delegate This Fully

The Structural Temptation

There is a structural temptation, particularly in technically complex domains, for boards to defer entirely to management. This instinct, while understandable, is misaligned with the nature of AI risk.

Why Deferral Fails

AI governance failures are not operational failures in the conventional sense — they are governance failures. They arise from:

- The absence of a defined risk appetite
- Unchecked automation bias
- AI systems that were never mapped
- Public claims that were never verified

> These are precisely the conditions that fiduciary responsibility is designed to prevent. The board does not need to understand the mathematics of a machine learning model. It needs to understand whether the organisation has a governance framework commensurate with the risk it is carrying.

Risk Taxonomy

AI risk is not a monolithic category. It comprises six distinct risk dimensions, each with its own governance implications — and each capable of amplifying the others.

Risk Category	Governance Implication
Strategic Risk	AI investments misaligned with enterprise strategy, or risk appetite not defined — leading to unchecked automation of consequential decisions.
Operational Risk	Model drift, data quality failures, hallucinations in customer-facing systems, and inadequate human-override mechanisms create process failures at scale.
Regulatory Risk	Non-compliance with DPDPA automated decision-making obligations, RBI IT governance expectations, SEBI algorithmic accountability requirements, and equivalents under DORA or the EU AI Act.
Cyber & Data Risk	GenAI-specific threat vectors, including prompt injection, data poisoning, and excessive agency in agentic AI systems, are compounded by expanded attack surfaces from third-party AI integrations.
Reputational Risk	AI washing through unsubstantiated public claims; consumer harm from AI outputs (as in the Air Canada precedent); bias in high-stakes decisions such as credit, hiring, and insurance pricing.
Capital Allocation Risk	AI spending without commensurate governance, leading to write-downs, regulatory provisioning requirements, and erosion of investor confidence.

These risks do not operate in isolation. Operational risk (a model hallucinating in a customer-facing application) generates reputational risk, which in turn generates regulatory risk, which in turn affects capital access and enterprise valuation. The interconnected nature of AI risk requires governance structures that cut across committee silos — which is precisely why board-level coordination is indispensable.

Regulatory Landscape — India

India's regulatory environment for AI governance is evolving from principle to enforcement. Boards would be well-advised to treat the following not as future obligations, but as current fiduciary standards:

DPDPA — Automated Decision-Making

The DPDPA creates obligations for data fiduciaries around the automated processing of personal data. As regulations are operationalised, obligations around automated decision-making affecting individuals will require documented accountability frameworks — directly implicating board oversight.

RBI IT Governance & Outsourcing Circulars

RBI's IT governance and outsourcing circulars extend risk management expectations to technology decisions, including those involving AI vendors and embedded AI features in licensed software. Third-party AI risk is squarely within the board's oversight perimeter.

SEBI — Algorithmic Accountability

SEBI's regulatory attention to algorithmic risk in capital markets contexts signals that explainability, auditability, and human override — the core tenets of AI governance — are becoming regulatory expectations, not merely best practice.

Sectoral Regulators — IRDAI & PFRDA

Sectoral regulators, including IRDAI and PFRDA, are similarly expanding their digital governance expectations in ways that will increasingly capture AI-enabled processes.

Global Parallels & Board Fiduciary Responsibility

EU AI Act

Classifies AI systems in high-stakes domains — including credit scoring, employment, and critical infrastructure — as 'high-risk,' requiring conformity assessments, human oversight mechanisms, and transparency obligations. Indian enterprises with European operations or investor exposure face direct obligations.

DORA

Imposes operational resilience requirements on financial entities that encompass AI-driven processes and third-party technology risk — a standard that is influencing regulatory thinking in multiple jurisdictions.

SEC AI Washing Enforcement

Charging investment advisers for false or misleading statements about AI capabilities has established a global disclosure standard. The principle is clear: AI claims must be truthful, evidentiary, and controlled.

Board Fiduciary Responsibility

> The convergence of these regulatory frameworks creates a clear fiduciary standard for boards: the absence of structured AI governance is itself a governance failure. Directors who cannot demonstrate that they have asked the right questions, reviewed appropriate management information, and exercised informed judgment over material AI risk may face personal liability exposure in the event of a significant AI-related failure.

This is not a theoretical concern. Regulatory frameworks explicitly contemplate senior management accountability. Board minutes, committee reports, and management information frameworks will constitute evidence — or its absence — in the event of investigation.

Strategic & Capital Implications

Enterprise Value

AI governance is rapidly becoming a component of enterprise value assessment. PE and institutional investors are beginning to embed AI governance due diligence into their pre-investment frameworks. The question is no longer merely 'do you use AI?' but 'can you demonstrate that you govern it?' Enterprises that cannot answer the second question with evidence face valuation discounts that are likely to widen as regulatory frameworks mature.

Capital Provisioning

Unresolved AI risk translates into capital exposure. Regulatory penalties, litigation costs arising from AI-related consumer harm, and the cost of emergency remediation following an AI-related incident can be material. More significantly, regulatory bodies are beginning to consider whether institutions with inadequate AI governance should be subject to additional capital requirements.

Competitive Positioning

There is a paradox at the heart of AI governance: the enterprises that invest in it most rigorously are likely to deploy AI most sustainably. Ungoverned AI creates operational fragility. Governed AI, by contrast, creates institutional confidence that enables more ambitious, higher-value deployment.

Investor Expectations

The ESG investment community is beginning to incorporate AI ethics and governance into its assessment frameworks. What began as a reputational consideration is becoming a measurable criterion. The 'G' in ESG — governance — is the natural home for AI accountability. Boards that frame AI governance within their ESG reporting will be better positioned to meet the growing expectations of institutional investors.

Operating Model Considerations

Governance Structures

No single board committee can carry the full burden of AI governance. The appropriate model assigns joint ownership across committees, with the full board receiving periodic integrated reporting:

Risk Committee

AI risk appetite definition, high-impact system classification, incident reporting, and third-party AI risk oversight.

Audit Committee

Internal audit coverage of material AI systems, independent model validation frameworks, and controls assurance.

Nomination & Governance Committee

Board AI literacy, director education programmes, and governance framework periodic review.

ESG Committee

AI ethics, algorithmic bias in customer-facing and employment decisions, and AI governance disclosure in sustainability reporting.

Management Reporting Frameworks

A single quarterly AI Governance Dashboard, presented to the full board or relevant committee, should consolidate the following metrics:

- AI inventory coverage — percentage of business units and products mapped to a documented AI register
- High-impact system classification status — percentage of consequential AI systems with documented risk assessments and human-override controls
- AI incident log — hallucinations, data leakage events, bias flags, and mean time to containment
- Compliance posture — regulatory obligations mapped and evidenced
- AI claims audit — percentage of external AI claims backed by evidence packs and legal review

Talent & Capability

AI governance requires a distinct capability set that sits between technology and risk management. Chief Risk Officers and Chief Compliance Officers need to understand AI-specific risk concepts — prompt injection, model drift, automation bias, agentic AI permissions — at a sufficient level to exercise meaningful oversight. Boards should satisfy themselves that management has invested in this capability, not merely in AI deployment.

Ecosystem & Vendor Risk

A significant and frequently underestimated dimension of AI governance is the risk embedded in 'vendor AI' and 'embedded AI' — AI features within licensed software platforms, cloud services, and third-party APIs. Many organisations have substantial AI exposure through their technology stacks, without a clear inventory or risk assessment of these embedded capabilities. This is a specific governance gap that boards should probe.

AI Governance Maturity Framework

The following five-level framework provides boards with a structured lens for assessing their organisation's AI governance posture. It is intended as a diagnostic tool, not a benchmarking exercise.

Level	Posture	Characteristics
Level 1	**Reactive**	No AI register. AI deployments are approved ad hoc — incidents managed after the fact. Board receives no AI-specific reporting — governance absent in substance, if not in form.
Level 2	**Structured**	A partial AI inventory exists. Basic risk classification underway. Audit Committee is aware of material systems. Incident reporting is manual and retrospective. Governance is emerging but fragmented.
Level 3	**Integrated**	Complete AI register maintained. High-impact systems are classified and risk-assessed. Human-override controls documented. Quarterly board reporting is in place. Internal audit coverage of material AI systems established.
Level 4	**Strategic**	AI risk appetite formally defined and board-approved. AI governance is embedded in capital allocation decisions. Independent model validation for material systems. External AI claims are subject to legal review. Vendor AI is included in the governance perimeter.
Level 5	**Embedded Governance**	AI governance is indistinguishable from enterprise risk governance. Real-time monitoring. Continuous assurance. AI ethics embedded in ESG reporting. Board AI literacy programme is active. Organisation can evidence governance to regulators and investors on demand.

Most Indian enterprises currently sit between Level 1 and Level 2. The regulatory and investor trajectory strongly suggests that Level 3 will become the de facto minimum expectation within the next 24 months. Boards that wait for regulatory compulsion to reach Level 3 will find themselves managing a crisis rather than building capability.

Board-Level Questions to Ask Management

These twelve questions should be treated as standing agenda items — not one-time enquiries. They are informed by OECD accountability principles and NIST's GOVERN function. Each question is designed to surface evidence, not merely assurance.

Strategy & Risk Appetite

1

Which specific business outcomes will AI materially change in the next 12 to 24 months — and what is our formally defined AI risk appetite, including which decisions we have determined not to automate?

2

What is the total AI spend in this financial year — and what governance approval was required before each significant AI investment was approved?

Inventory & Classification

1

Do we have a complete and current AI register — including AI features embedded in licensed vendor software, cloud platforms, and third-party APIs — and when was it last independently verified?

2

Which AI systems have been classified as 'high-impact' — specifically those involved in credit decisions, hiring, pricing, fraud flags, and KYC — and under what classification framework?

Human Oversight & Accountability

1

For every high-impact AI system, can management demonstrate where and how a human can override, stop, or reverse the AI's output — and is automation bias actively monitored and reported?

2

Who is the named accountable senior executive for each material AI system — and how is that accountability documented and reviewed?

Security & Resilience

1

Have we tested for GenAI-specific security threats — including prompt injection, data poisoning, and excessive agency in agentic AI systems — and what are the current open findings?

Assurance & Audit

1

What internal audit and independent validation coverage exists for our material AI systems — and has the Audit Committee reviewed the findings?

2

What is our AI incident rate, and what is the mean time to containment for AI-related issues, including hallucinations, data leakage, and bias events?

Disclosure & Regulatory Readiness

1

Are all public statements about our AI capabilities — in investor communications, marketing materials, and regulatory submissions — truthful, evidentiary, and subject to legal review before publication?

2

Has management prepared a regulatory readiness assessment against DPDPA automated decision-making obligations, RBI IT governance requirements, and any relevant global frameworks applicable to our business — and is the board satisfied with the gap closure plan?

Case Illustrations

Case A: The Disclosure Liability — 'AI Washing' in Financial Services

Consider a financial services firm that, during a period of intense competitive pressure, begins describing its credit underwriting process as 'AI-powered' in investor presentations and marketing materials. The reality is that a basic rule-based scoring model, with a small ML component added eighteen months earlier, drives the majority of decisions. The language is not fabricated — but it is materially misleading.

This scenario mirrors precisely the SEC's enforcement action against two investment advisers for false or misleading statements about their AI use. The governance failure is not in the technology — it is in the absence of a disclosure review process that subjects AI claims to the same evidentiary standard as financial statements. An Indian equivalent is a foreseeable risk as SEBI's regulatory attention intensifies. The remedy is simple in concept: every external AI claim should be backed by an evidence pack and reviewed by legal counsel before publication. The board should ask whether this process exists.

Case B: The Accountability Gap — Customer-Facing GenAI

A large retail bank deploys a GenAI-powered customer service interface. The system is capable, fast, and dramatically reduces call centre volumes. Eighteen months into deployment, a customer files a complaint having acted on incorrect information provided by the chatbot regarding a loan prepayment waiver — information the human customer service team would not have provided.

This scenario aligns with the Air Canada case, in which a court found the airline liable for negligent misrepresentation after its chatbot provided inaccurate policy information. The finding was unambiguous: the organisation is accountable for AI outputs presented to customers. The governance failure was not technical — the model had known limitations. The failure stemmed from the absence of a human escalation protocol, inadequate disclaimer architecture, and a lack of board-level visibility into the risk. The remedy requires documented human-override mechanisms, customer-facing disclosure, and board awareness of the risk exposure.

Case C: The Silent Inventory — Vendor AI Risk

An insurance company's board approves a technology transformation programme. The programme deploys a cloud-based policy administration system, a CRM platform, and a third-party analytics solution. What the board is not told — because no one has assessed it — is that each of these platforms includes embedded AI features: automated document classification, AI-driven customer segmentation, and predictive churn modelling. None of these is in the AI register, because no AI register exists.

This is not an exceptional scenario. It is the norm in most mid- to large-sized Indian enterprises. The governance implication is significant: the organisation is carrying AI risk that it has not classified, assessed, or reported. When a regulatory enquiry or incident exposes this gap, the board's position is considerably weaker than it need have been. The remedy is a comprehensive AI inventory, including vendor AI and embedded features, as a prerequisite for any meaningful AI governance programme.

Case D: The Governance Leader — Proactive Architecture

Contrast the above with a financial services group that, three years ago, established an AI Governance Committee as a sub-committee of the Risk Committee, with independent director participation. The committee commissioned a full AI register, classified high-impact systems, required human-override documentation as a condition of production deployment, and established a quarterly AI dashboard presented to the full board.

When a GenAI feature in a third-party KYC platform produced anomalous outputs that affected a cohort of customer verifications, the incident was identified within hours, contained within days, and reported to the regulator proactively within the required timeframe. The board was informed through an existing reporting channel, not through a crisis briefing. The regulatory response was measured. The reputational impact was managed. The governance infrastructure, built before the crisis, proved its value precisely when it was needed.

12–24 Month Action Roadmap

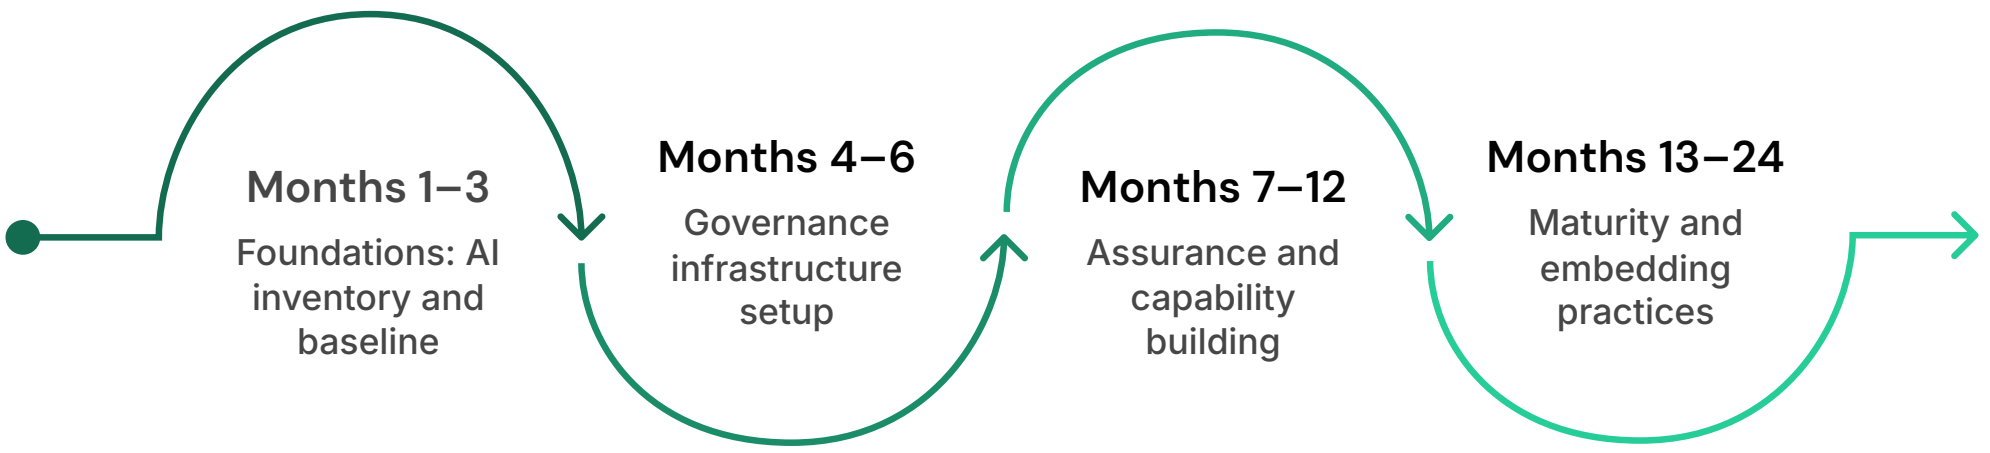

Months 1–3: Foundations

- Commission a comprehensive AI inventory covering all production AI systems, vendor AI, and embedded AI features across business units. This is the governance prerequisite for everything that follows.
- Define the board's AI risk appetite — explicitly, including which decisions will not be automated — and have this formally approved at a board meeting with the decision minuted.
- Assign committee ownership: Risk Committee for risk appetite and classification, Audit Committee for assurance and controls, Nomination Committee for board literacy, ESG Committee for ethics and disclosure.
- Conduct a disclosure audit: review all external AI claims in investor communications, marketing, and regulatory submissions against an evidentiary standard.

Months 4–6: Governance Infrastructure

- Classify all AI systems by impact level. Designate 'high-impact' systems and document human-override mechanisms, accountable owners, and incident escalation protocols for each.
- Establish the quarterly AI Governance Dashboard. Define the metrics: inventory coverage, incident rate, mean time to containment, compliance posture, and claims audit status.
- Engage internal audit to scope coverage of material AI systems. Establish the assurance framework.
- Conduct a regulatory readiness assessment against DPDPA, RBI IT governance requirements, and any applicable global frameworks.

Months 7–12: Assurance & Capability

- Commission independent model validation for high-impact AI systems. Present findings to the Audit Committee.
- Test for GenAI-specific security threats — prompt injection, data poisoning, excessive agency — and review open findings at the Risk Committee.
- Initiate a board AI literacy programme. Directors do not need technical expertise; they need sufficient conceptual grounding to ask the right questions and assess the quality of management responses.
- Begin mapping to ISO/IEC 42001 as the management system backbone for AI governance accountability and continuous improvement.

Months 13–24: Maturity & Embedding

- Integrate AI governance into capital allocation processes. Require an AI governance assessment as a prerequisite for approval of significant AI investments.
- Embed AI ethics and governance disclosure within ESG reporting. This positions the enterprise to engage with institutional investors and to anticipate regulatory reporting requirements.
- Conduct an annual board-level AI governance review — not a management briefing, but a structured board discussion informed by the full year's dashboard data, audit findings, and regulatory developments.
- Target Level 3 (Integrated) maturity on the framework presented in this Board Brief. For regulated financial institutions, Level 4 (Strategic) should be the 24-month objective.

Strategic Conclusion

There is a version of this conversation that boards have been deferring — in which AI governance is treated as a management matter, a technology item, or a regulatory horizon to be monitored. That version of the conversation is no longer available.

The regulatory frameworks are converging. The enforcement actions have begun. The consumer harm precedents have been set. The investor due diligence frameworks are being built. The twelve questions in this Board Brief are not theoretical. They are the questions that will be asked — by regulators, by investors, by courts — in the event of a significant AI-related failure. The only question is whether the board was asking them first.

> What distinguishes genuine board leadership in this domain is not technical knowledge — it is the intellectual discipline to demand evidence, to define appetite, to assign accountability, and to require assurance. These are governance skills. They are the skills boards already exercise in managing financial, credit, and operational risk. AI governance requires the same rigour applied to a new and consequential domain.

A Closing Reflection

> The enterprises that govern AI well will not merely avoid failures. They will build the institutional confidence to deploy AI more ambitiously, more sustainably, and to greater competitive effect than those who mistake speed for advantage. Governance is not a brake on AI adoption. It is the infrastructure that enables ambitious AI adoption.

> The board's role is not to slow the organisation down. It is to ensure that when the organisation moves fast, it moves on foundations that can bear the weight.

APPENDIX A | BOARD-READY

One-Page Whitepaper Executive Summary

Suitable for Downloadable PDF Distribution

Aparna Kumar | Founder — Nexora Tech I Former CIO — SBI & HSBC | Board Advisor |

Why This Matters Now

- AI governance failures are governance failures, not technology failures — and they carry board-level fiduciary consequences.
- Regulatory exposure is active: DPDPA, RBI IT frameworks, SEBI algorithmic accountability, EU AI Act, DORA, and SEC AI washing enforcement.
- Consumer harm precedents are set (Air Canada chatbot liability). Indian enterprises face equivalent exposure.
- Investor due diligence frameworks are incorporating AI governance as a valuation criterion.

The Six Risk Dimensions

Strategic | Operational | Regulatory | Cyber & Data | Reputational | Capital Allocation

These risks are interconnected and amplifying — operational failures generate regulatory and reputational cascades.

Governance Framework Anchors

- **NIST AI RMF:** Govern / Map / Measure / Manage — the board's risk language
- **ISO/IEC 42001:** Certifiable management-system backbone for AI accountability
- **OECD AI Principles:** Fairness, transparency, robustness, accountability — publicly adoptable norms

Immediate Board Priorities

01

Commission a complete AI register — including vendor AI and embedded AI features

02

Define and formally approve the AI risk appetite at the board level

03

Assign committee ownership: Risk, Audit, Nomination, ESG

04

Audit all external AI claims for truthfulness and evidentiary basis

05

Establish the quarterly AI Governance Dashboard

Metrics Boards Should Require

AI Inventory Coverage

% of business units mapped to the AI register

High-Impact Systems

% with documented risk assessments and human-override controls

AI Incident Rate

Mean time to containment for AI-related issues

AI Claims Audit

% of external AI claims backed by evidence packs and legal review

Target Maturity

1

Level 1–2

Most Indian enterprises currently: Reactive to Structured

2

Level 3

Regulatory minimum expectation within 24 months: Integrated

3

Level 4

Regulated financial institutions — 24-month objective: Strategic

Thank You

We deeply appreciate your engagement with this Board Brief. Your commitment to rigorous AI governance is paramount in shaping a future where technological innovation thrives responsibly.

Our mission at Nexora Tech is to empower boards with the insights and frameworks necessary to navigate the complexities of AI governance, transforming potential risks into strategic opportunities and driving sustainable value creation.

We believe that effective governance is not merely about compliance, but about building the institutional confidence required for ambitious and ethical AI deployment. Thank you for being a leader in this critical domain.

Clarity. Architecture. Governance. Impact.

Aparna Kumar | Founder — Nexora Tech I Former CIO — SBI & HSBC | Board Advisor |

Email: sales@nexoratechsolutions.com

WhatsApp: +91 8799918172

www.nexoratechsolutions.com