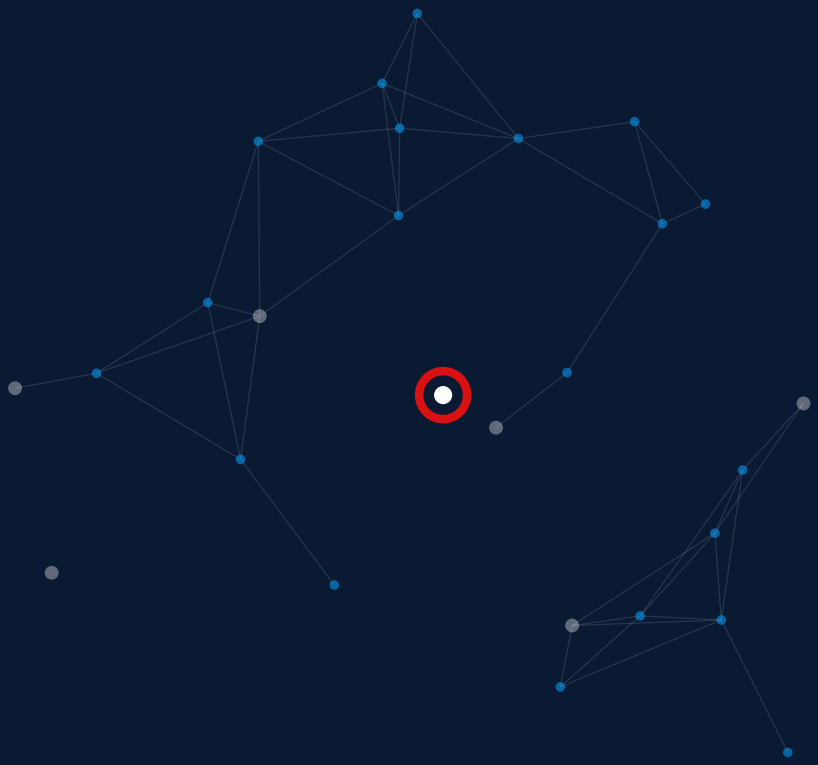
KRYNTHNODE



AI and technology strategy | Brief

Generative AI governance

A board-level brief

Published 1 April 2026
Krynthnode Research and Advisory LLP | krynthnode.com

Research that shows its working.

Executive summary

Generative AI is being adopted faster than most organisations can govern it. Employees use public tools, vendors embed AI features into existing software, and business units build their own applications. Boards are increasingly asked, by regulators, investors and customers, how they oversee these risks.

Directors do not need to be technical experts to oversee AI well. They need clear accountability, an accurate picture of where AI is used, and assurance that the most significant risks are managed. This brief sets out the questions a board can reasonably ask management, what good answers look like, and a one-page checklist, drawing on the NIST AI Risk Management Framework and ISO/IEC 42001.

Key findings	Confidence
Clear executive accountability is the foundation of AI oversight. Boards should expect one named executive to be answerable for AI risk across the organisation.	High confidence
An inventory of AI systems, including AI features embedded in vendor products, is a prerequisite for any credible oversight.	High confidence
Incident response plans rarely cover AI-specific failure modes such as harmful outputs, prompt injection or data leakage through AI tools.	Moderate confidence
Recognised frameworks such as the NIST AI RMF and ISO/IEC 42001 give boards a structure for oversight without requiring technical depth.	High confidence

Each finding is graded in the Krynthnode confidence ledger. See the methodology at the end of this report.

What we recommend

1. Name one executive accountable for AI risk and set a regular reporting rhythm to the board or its risk committee.
2. Require an inventory of AI systems in use, classified by risk.
3. Approve an AI use policy covering data, human oversight, testing and acceptable use of public tools.
4. Ask for AI-specific scenarios in incident response plans and exercises.
5. Review how AI use affects obligations under existing law, including the DPDP Act, and in overseas markets such as the EU.

© 2026 Krynthnode Research and Advisory LLP. All rights reserved.

Why boards are being asked

AI changes the risk profile of an organisation in ways that fall squarely within a board's oversight duties. Generative AI can produce inaccurate or harmful content, expose confidential or personal data, infringe intellectual property, embed bias in decisions and be manipulated by attackers. At the same time, failing to adopt AI can carry strategic risk.

Regulation is also developing. The European Union's AI Act, which entered into force in August 2024, imposes obligations that can apply to Indian companies serving European markets. In India, organisations must apply existing laws, including the Information Technology Act and the Digital Personal Data Protection Act, to their use of AI, and should follow government guidance as it develops.

Ten questions for management

Question	What a good answer looks like
Who is accountable for AI risk?	One named executive, with authority across business units
Where do we use AI today?	A current inventory, including vendor-embedded AI, classified by risk
What are our highest-risk uses?	Those affecting customers, credit, health, safety or legal rights, with named owners
What rules govern AI use?	An approved policy covering data, oversight, testing and public tools
How is personal data protected?	AI use assessed against DPDP obligations; no personal data in unapproved tools
How do we test before deployment?	Documented testing for accuracy, bias, security and misuse
Who can override AI decisions?	Human oversight defined for consequential decisions
What happens when AI fails?	Incident plans with AI-specific scenarios, rehearsed
How do we manage AI vendors?	Contract terms on data use, model changes and incident notification
How do we measure value?	Business outcomes tracked against the investment

© 2026 Krynthnode Research and Advisory LLP. All rights reserved.

Frameworks that help

Boards do not need to invent a governance model. The NIST AI Risk Management Framework organises AI risk management into four functions: govern, map, measure and manage. Its Generative AI Profile, published in July 2024, addresses risks specific to generative systems. ISO/IEC 42001, published in December 2023, is the first international standard for an AI management system; organisations can adopt it as a framework or seek certification.

For security risks specific to applications built on large language models, the OWASP Top 10 for Large Language Model Applications is a widely used reference. It lists prompt injection, in which instructions hidden in content cause a model to behave in unintended ways, as the leading risk.

The one-page oversight checklist

- An executive is named as accountable for AI risk, and reports at least twice a year.
- An AI inventory exists, is current and includes AI embedded in vendor products.
- High-risk AI uses are identified, with owners and controls.
- An AI use policy is approved and communicated to all staff.
- Personal data use in AI has been assessed against the DPDP Act.
- Testing standards exist for AI systems before deployment.
- Incident response plans include AI failure scenarios and have been exercised.
- Vendor contracts address data use, model changes and incident notification.
- The board has received education on AI opportunities and risks in the last year.

Methodology and the confidence ledger

This brief draws on the frameworks, standards and regulations listed in the references, applied to the governance responsibilities of boards. It is not legal advice.

Grade	What it means	How to use it
High confidence	Consistent across several independent sources and established practice.	Suitable for committing budget.
Moderate confidence	Supported, with some gaps or variation between contexts.	Suitable for a pilot with a review date.
Emerging evidence	Limited or recent evidence.	Worth monitoring before acting.

References

1. National Institute of Standards and Technology. AI Risk Management Framework (AI RMF 1.0). January 2023.
2. National Institute of Standards and Technology. NIST AI 600-1, Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile. July 2024.
3. International Organization for Standardization. ISO/IEC 42001:2023, Artificial intelligence: Management system. December 2023.
4. European Union. Regulation (EU) 2024/1689 (Artificial Intelligence Act).
5. OWASP Foundation. OWASP Top 10 for Large Language Model Applications, 2025 edition.

© 2026 Krynthnode Research and Advisory LLP. All rights reserved.

About Krynthnode

Krynthnode Research and Advisory LLP is an independent technology and cybersecurity research and advisory firm. We help leaders make decisions they can defend: every finding we publish carries its evidence, its sources and an honest rating of how sure we are. We do not sell technology, take referral fees or accept payment for ratings.

To discuss this research with the analyst team, visit krynthnode.com/contact or email hello@krynthnode.com.

This publication reflects the views of Krynthnode analysts based on the sources available at the time of writing. It is provided for information and is not legal, regulatory, financial or investment advice. Mention of any organisation, product or standard is not an endorsement. Trademarks belong to their respective owners. No part of this publication may be reproduced or redistributed without written permission, except for short quotations with attribution.

© 2026 Krynthnode Research and Advisory LLP. All rights reserved.