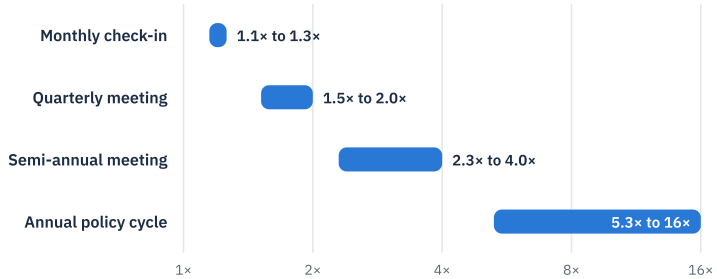
Is Your Board Keeping Up With AI?

Chris B. Kennedy · A call to action for boards of large enterprises, public companies and regulated industries

AI changes more between two board meetings than governance was built to absorb, and published guidance has not caught up. Boards need questions and checklists they can use at the next meeting.

1 How much AI changes between your meetings



Capability growth between meetings (log scale), at a doubling every 3 to 5 months

AI capability is doubling roughly every 3 to 5 months. A board that meets quarterly governs a picture that is already half out of date; an annual policy can be sixteen times behind.

The binding constraint is no longer the technology. It is how fast the organization, and its board, can absorb change safely.

The governance gap, and what agile oversight looks like: “AI Moves Faster Than Your Board Meets.”

2 The guidance has not caught up

Of 34 director bodies, consultancies, standards setters and regulators benchmarked, how many ask:

Does our oversight cadence keep pace with AI?	4
What can agents do alone, and can we stop them?	5
What P&L impact has AI actually delivered?	5
Who is accountable for AI-produced work?	2
Does AI doing a person's job get their controls?	1
Are our people keeping real expertise?	1
What are we deliberately not funding?	0

For comparison, 16 of 34 ask about identifying and managing AI risk. The FSB's October 2026 sound practices never mention return on investment and call agent practices "nascent."

3 Top AI risks to put on the agenda

RISK	ORGS
Vendor and model concentration. A few providers change terms or price, or fail.	18
Data, privacy and IP leakage. Client or proprietary data goes into AI tools.	13
Errors that reach customers. Wrong or biased output no one catches.	13
Regulatory exposure. New AI rules and existing law applied to AI.	10
AI-enabled fraud. Cloned voices, fake documents, automated phishing.	6
Agents acting beyond authority. AI that moves money or changes systems unchecked.	5

Orgs: outside organizations, of 34, raising the risk. Agent risk is growing quickly and is still rarely asked.

4 Any framework, plus what AI changes

Anchor AI in the enterprise risk framework you already run and the NIST AI Risk Management Framework (for financial institutions, the FSB's twelve sound practices), then add three things:

- **Human ownership of every control.** Every AI control, human or technical, has a named owner.
- **Velocity oversight.** Oversight, talent and spend, cyber and data controls that move at AI speed.
- **Exponential complexity oversight.** Purpose-built agents, limited to what can still be modeled and stopped.

5 Before your next meeting

1. **Ask the ten questions.** Find where answers are thin.
2. **See all the AI.** An inventory including vendor-embedded and employee tools, each with an owner.
3. **Set escalation triggers.** Agents that move money, new sensitive data, a vendor change or an incident reach the board at once.
4. **Get one page every meeting.** What changed and what is new on the risk list.
5. **Judge the answers.** Mark each one clear, partial or missing; close the gaps within two meetings.

The full Board AI Toolkit

→ Board AI Checklist (15 items) and framework → Independent review and scorecard → Board and leadership training
→ AI advisory and monitoring → Sample policies, evaluations and evidence guides → Talent assessment and AI hiring

Available through an engagement. See What's Next.