

Board White Paper #3

Risk, Regulation & Trust

A board-level perspective on governing AI risk, regulatory exposure, and institutional trust as AI systems scale

Prepared by Dr. Ori Marom
 Managing Director, Segmentis B.V.

Board purpose	To support board discussion on how AI-related risks should be identified, governed, and translated into credible oversight as regulation and public scrutiny increase.
Primary board question	How can the organization scale AI while maintaining regulatory compliance, clear accountability, and institutional trust?
Intended audience	Supervisory board members, non-executive directors, executive board members, and senior leaders responsible for AI governance, risk, compliance, and transformation.

Executive summary

AI risk has moved from a technical control issue to a board-level question of institutional trust, regulatory accountability, and strategic resilience. As AI systems become embedded in decisions that affect customers, employees, markets, and public confidence, boards must ensure that governance keeps pace with both deployment and scrutiny.

The EU AI Act makes this shift explicit by introducing a risk-based framework that requires organizations to understand where AI is used, what impact it may have, who is accountable, and how oversight can be demonstrated. For boards, the practical challenge is not only formal compliance, but the ability to show that AI-supported decisions remain explainable, proportionate, and subject to meaningful human control.

The paper argues that effective AI governance must integrate risk, regulation, and trust into execution. Compliance mechanisms should not sit outside the business as static documentation. They should become operational guardrails that help teams move faster with clearer accountability, stronger escalation paths, and continuous monitoring as AI systems evolve.

1. Board context and relevance

This white paper addresses the board’s responsibility to ensure that AI-related risk is governed in a way that protects institutional trust, regulatory standing, and long-term accountability. It focuses on how boards should understand and oversee AI systems in an environment where regulation, public scrutiny, and societal expectations are rapidly converging. In Europe, this convergence is now embodied in the EU AI Act³.

This shift has made AI risk more consequential for boards. The EU AI Act is not merely a regulatory development, but a signal that AI systems are increasingly being used in contexts where errors, bias, or loss of accountability can have material consequences. For boards, the implication is that AI risk can no longer be treated as a technical concern alone. It has become a strategic and fiduciary issue.

The practical challenge is no longer whether AI can be governed in principle, but whether organizations can assign responsibility, maintain trust, and remain compliant as AI systems scale faster than traditional governance mechanisms. Boards that fail to address this gap risk creating exposure precisely where AI becomes most material to customers, employees, markets, and regulators^{1,2,5,6}.

2. Board questions for early discussion

Boards engaging with AI when subject to the EU AI Act often ask:

1. Are we confident that our AI-related risks are identified and managed with the same discipline as other regulated risks?
2. Which uses of AI in our organization fall into higher-risk categories under the Act: *Unacceptable-Risk* and *High-Risk*? And why?
3. Where does our existing risk and compliance framework still work, and where does it need to evolve?
4. How do we remain compliant as AI systems change, learn, and scale over time?
5. Which of our actual or intended AI uses could undermine trust and/or ethical standards; even if we remain formally compliant?

The purpose of these questions is not to slow AI adoption, but to ensure that responsibility, compliance, and trust are clear before AI systems become too embedded to govern effectively.

3. Risk-based governance in practice

AI brings both material opportunity and material risk. In a recent survey by Moody's, 96% of respondents said that risk and compliance roles will change, becoming more strategic, more focused on oversight, and more dependent on a deeper understanding of AI systems⁴.

The EU AI Act reflects a broader recognition among policymakers that AI risks are not hypothetical. Financial and technology reporting increasingly points to situations in which AI systems were operational long before accountability had been clearly assigned, resulting in public backlash and regulatory intervention.

At its core, the Act introduces a **risk-based approach**. AI systems are classified according to their potential impact, with stricter obligations where AI affects safety, fundamental rights, or access to essential services. The Act does not prohibit innovation. It requires **clarity of responsibility, transparency of use, and the ability to intervene**.

For boards, the legal detail matters less than the underlying logic: **AI risk scales with impact, and so must governance**. Compliance should therefore not be treated as a one-off exercise. It should be embedded in the way AI-supported decisions are designed, deployed, and monitored.

Trust is the missing link. Organizations often discover that trust erodes before formal regulatory limits are breached. The EU AI Act responds to this reality by requiring organizations to demonstrate not only that AI systems function, but that their use is defensible and accountable.

Segmentis' **Minimum Viable AI Governance (MV-AIG)** is relevant here not as a stand-alone compliance tool, but as a way to ensure that governance keeps pace with evolving risk. It provides the minimum structure required to demonstrate responsibility when AI systems come under scrutiny.

4. How proper understanding of risks informs AI governance

Risk understanding is what turns regulatory obligation into effective governance.

In **White Paper #2**, we argued that governance must follow decisions. The EU AI Act reinforces the same principle by asking, implicitly but forcefully: *Who is responsible for this AI-supported outcome? And: Can that responsibility be demonstrated?*

Where organizations often struggle is not in understanding the Act itself, but in translating its risk-based logic into governance practice. Overly generic controls create friction without protection. Insufficient controls in high-impact areas create exposure that no compliance function can repair later^{6,7,8}.

Segmentis helps boards align governance structures with actual risk exposure. High-risk AI uses require explicit ownership, human oversight, and escalation mechanisms. Lower-risk uses require proportionate controls that preserve agility.

When governance reflects real risk, compliance becomes sustainable rather than reactive.

5. Designing compliance mechanisms that enable execution

Compliance mechanisms fail when they are detached from execution.

Many organizations experience the EU AI Act as a threat to speed and innovation. This perception is often misguided.

Recent business reporting shows that enforcement and reputational damage tend to arise not from deliberate non-compliance, but from fragmented responsibility and unclear operational boundaries^{5,6,8}. Compliance processes that sit outside execution slow organizations down without reducing risk.

Effective boards insist that compliance mechanisms are *execution aware*. Regulatory requirements are translated into **operational guardrails** that teams understand and can work within. Responsibilities are assigned upfront. Monitoring is continuous, not episodic.

This prepares the organization for **execution reality**, addressed in White Paper #4. AI initiatives fail not because regulation exists, but because organizations cannot reconcile compliance with day-to-day decision-making.

Boards that integrate regulation, risk, and trust into execution design move faster and with greater confidence. The EU AI Act then becomes a forcing function for clarity, not a brake on innovation.

6. Indicators that the board is on the right path

Boards know they have struck the right balance between execution speed, regulatory compliance, and institutional trust when AI governance becomes anticipatory rather than defensive. Regulatory

engagement is no longer treated as a reaction to external pressure, but as part of how the organization understands and manages the consequences of AI-supported decisions.

Evidence of progress appears when AI risks are surfaced early, ethical concerns are discussed openly, and trust is maintained even when systems come under scrutiny. These are signs that governance is not merely documenting compliance, but shaping the conditions under which AI can be used responsibly.

Misalignment is equally visible. Novel risks that are ignored eventually become incidents, and governance that lags deployment gradually erodes credibility. At that point, the problem is no longer only whether the organization is compliant, but whether it can still defend the way AI is being used.

Boards that integrate risk, regulation, and trust into ordinary governance create the conditions for AI to scale responsibly. At that point, AI oversight becomes less about reassurance after the fact and more about disciplined confidence before decisions are made.

7. Board implications

Board action	Implication for AI risk, regulation, and trust
Map AI exposure	Boards should require a clear view of where AI is used, which systems may fall into higher-risk categories, and where regulatory obligations are likely to be material.
Assign accountability	High-impact AI use should have explicit ownership across business, technology, risk, and compliance so responsibility can be demonstrated when systems are scrutinized.
Embed compliance in execution	EU AI Act readiness should be translated into operational guardrails that teams can work within, rather than treated as documentation outside the business.
Protect trust explicitly	Boards should test whether AI-supported decisions remain explainable, proportionate, and defensible to customers, employees, regulators, and other stakeholders.
Monitor as systems evolve	Governance should remain continuous as AI systems change, learn, and scale, with clear escalation when risk exposure or public expectations shift.

Taken together, these actions help boards move AI oversight from periodic compliance reassurance to an ongoing discipline of accountable, trusted, and regulation-ready execution.

8. Case study – When EU AI Act compliance exposed a trust dilemma

A European financial services firm introduced AI to support parts of its credit assessment process. Early results were encouraging. Decisions were faster, model performance was strong, and initial internal reviews did not identify material control weaknesses.

The issue became visible when the organization began mapping the system against the **EU AI Act**. The credit model appeared likely to fall within the *High Risk* category because it influenced access to financial services. That classification changed the discussion. The firm now had to demonstrate clear documentation, data governance, human oversight, risk management, monitoring, and accountability for AI-supported decisions.

On paper, several of these obligations were partly covered. In practice, they were fragmented. Risk teams owned the compliance documentation, IT teams owned the model controls, and business units owned the customer outcomes. No single executive could credibly explain how the Act's requirements would be met end to end once the system was scaled.

The board also faced a **trust dilemma**. Disclosing the use of AI more openly could reassure regulators and strengthen transparency, but it might also unsettle customers who expected sensitive credit decisions to remain primarily human. Keeping the system technically compliant but poorly explained carried the opposite risk: the firm might satisfy formal obligations while weakening confidence if customers later discovered that AI had shaped lending decisions without a clear account of human review.

The board intervened before the gap became a regulatory or reputational incident. It asked management to treat EU AI Act readiness not as a *legal filing exercise*, but as an *operating-model issue*. Ownership for *High Risk* AI use was assigned explicitly, human oversight was clarified, customer-facing explanations were redesigned, and monitoring was connected to real credit decisions rather than to model documentation alone.

The result was not a slower innovation agenda. It was a more defensible one. The firm retained the benefits of AI in credit assessment while strengthening EU AI Act compliance, decision ownership, regulatory readiness, and trust in the way AI-supported decisions were made.

Selected references

1. The Economist (2025, September 22). *Why AI systems may never be secure, and what to do about it*.
2. The Economist (2025, September 25). *How to stop AI's "lethal trifecta"*.
3. European Union (2024). *Regulation (EU) 2024/1689 (Artificial Intelligence Act)*. Official Journal of the European Union, L 2024/1689.
4. Moody's (2025). *From reactive to proactive: how AI is transforming risk and compliance*.
5. OECD. (2026). *OECD due diligence guidance for responsible AI*. OECD Publishing.
6. Strategic Risk Global (2026, March 19). *Regulatory convergence puts pressure on GRC leaders to rethink governance*.
7. UK Information Commissioner's Office (2023). *AI and data protection: risk management guidance*.
8. World Economic Forum. (2024). *Governance in the age of generative AI: A 360° approach for resilient policy and regulation*. World Economic Forum White Paper.