

Who governs the AI in your organisation?

Why boards need visibility and evidence of AI-built and AI-driven applications

Edition	1 · October 2026
Reference date	8 October 2026
For	Board and executive team · IT, security and compliance
Legislation	GDPR · ePrivacy · AI Act · NIS2 (NL, BE, DE) · CRA · Data Act · PLD · DORA · EAA
Key message	DEMONSTRABILITY

1

Key figures

49–78%

of employees use AI tools their employer has not approved

51%

connect AI tools to work systems without IT approval

Personal

liability: directors in NL, BE and DE under NIS2 legislation

€35M / 7%

maximum AI Act fine (prohibited practices, worldwide turnover)

€20M / 4%

maximum GDPR fine

24 hours

first notice of a NIS2 incident and of an exploited vulnerability (CRA, since 11 Sept 2026)

Sources for these figures: see chapter 3 (research 2025–2026) and chapter 11.

Executive summary

AI is now part of almost every application your organisation uses: in tools built in-house, in purchased SaaS and in agents that carry out tasks on their own. Much of it has never passed through procurement, security or compliance.

Employees choose their own tools, connect them to business systems and build an application with AI in an afternoon. At the same time, the law has changed. Since NIS2 was transposed in the Netherlands, Belgium and Germany, cybersecurity has become a board responsibility with personal liability. The AI Act, the GDPR, the Cyber Resilience Act, the Data Act and the new product liability rules each add their own obligations on top.

KEY FINDING: all these laws ask for the same thing: **demonstrability**. Know which AI is running, who owns it, which data goes into it and what the system does. A policy on paper and a spreadsheet are not enough: applications change weekly, AI models are replaced and vendors quietly add AI features. Boards therefore need an AI governance tool that continuously registers, collects evidence and reports.

This whitepaper explains why, which laws apply and what such a tool must be able to do at a minimum. Part A is written for the board and executive team, part B for IT, security and compliance.

PART A · FOR THE BOARD AND EXECUTIVE TEAM

2 The incident

In autumn 2026, the author spoke with a senior manager at a listed ICT organisation operating in the Netherlands and Belgium. A commercial team there was using a signal-based go-to-market platform: software that collects buying signals from website visits, LinkedIn activity and job changes, and shows which individual decision-maker is “warm”.

The vendor's public information showed that:

- the platform delivers intent data at the **individual level**, about people who are not customers;
- it integrates with tools that make employees' LinkedIn accounts connect and send messages automatically, which LinkedIn prohibits in its user agreement;
- it uses AI to score individuals, without showing which models or sub-processors are involved.

Nobody had bad intentions. The team wanted to sell faster, and the tool did what it promised. But the organisation itself remained the controller under the GDPR, the deployer under the AI Act and, as an essential or important entity, responsible for its supply chain under NIS2.

The question the manager asked is the core of this whitepaper: **how do we, as a board, know which AI applications are running in our organisation, and can we prove they are allowed?**

The case has been anonymised. Neither the organisation nor the vendor is named.

3 The sprawl

The incident is not an exception. Recent research shows the same picture.

Finding	Source
49% of employees at organisations with 500+ staff use unapproved AI tools; 51% connect AI tools to work systems without IT	BlackFog, January 2026
52% of knowledge workers use unapproved AI; 58% of executives report an AI-related security incident in the past 12 months	Okta, AI Agents at Work 2026
78% of US employees use unapproved AI; only 7.5% received extensive AI training	WalkMe (SAP), August 2025
In IT and telecoms, 73% use shadow AI, while 84% know the risks	TrustedTech, September 2026

Three forms of sprawl

- 1. **Built in-house with AI.** Employees and teams build applications with AI coding assistants and no-code platforms. What used to be a months-long IT project now goes live in an afternoon, often without review, a test environment or an owner.
- 2. **Purchased SaaS with AI.** Existing vendors add AI features to software that has been in use for years. The contract, the data processing agreement and the risk assessment predate that AI.
- 3. **AI agents.** Systems that decide, send messages or process data on their own, with permissions across several applications at once.

What stands out is the gap between the board and the shop floor: executives believe they have visibility, while more than half of employees use their own tools. **That gap is exactly what makes directors vulnerable.**

4 Director liability in the EU

Cybersecurity is no longer an IT topic but a board duty, and that duty cannot be delegated. Article 20 of the NIS2 Directive requires every member state to ensure that management bodies approve security measures, oversee their implementation, follow training and can be held liable for infringements. Each EU country has transposed this into its own law.

Country	National law	What it means for directors
Netherlands	Cyberbeveiligingswet (Cbw), in force since 15 August 2026	Duty of care and training obligation for the board; personal liability for neglect; supervisors can, in certain circumstances, temporarily suspend a director
Belgium	NIS2 Act of 26 April 2024, in force since 18 October 2024	The management body approves the measures, oversees them and can be held liable; supervision by the CCB
Germany	NIS2 Implementation Act, in force since 6 December 2025	§ 38 BSIG: management must implement and monitor the measures, is personally liable for damage caused by breach of duty and must undergo training
Other EU	National NIS2 law	The same core from Article 20 NIS2, implemented nationally

On top of this, ordinary company law applies: improper performance of duties (Art. 2:9 BW in the Netherlands, Art. 2:56 WvV in Belgium, § 93 AktG and § 43 GmbHG in Germany). A director without visibility of AI applications that process personal data or access critical systems will find it hard to plead ignorance.

Three other laws add pressure

- **GDPR Art. 5(2).** The organisation must be able to demonstrate compliance; the burden of proof lies with it.

- **AI Act.** The deployer has its own obligations, such as AI literacy (Art. 4, since February 2025) and transparency (Art. 50, since August 2026).
- **Stock listing.** The annual report describes the principal risks and internal control. A structural gap in AI governance belongs there, and is read by supervisors, auditors and investors.

PART B · FOR IT, SECURITY AND COMPLIANCE

5

Heatmap: which law affects which AI application

Regulations apply directly and identically in all EU member states. Directives are transposed per country, so the name, supervisor and details differ. Below is an assessment of the impact per type of application.

Impact: ●●● high ●● medium ● low

Law	Applies	Built in-house	SaaS with AI	AI agents	Core for AI applications
GDPR	Regulation, since 2018	●●●	●●●	●●●	Legal basis, profiling, DPIA, data processing agreement, transfers outside the EEA
ePrivacy	Directive, transposed nationally	●●	●●●	●●	Consent for tracking and visitor identification
AI Act	Regulation; Art. 4 since Feb 2025, Art. 50 since Aug 2026, high-risk from Dec 2027	●●●	●●	●●●	Role (provider or deployer), AI literacy, transparency, risk class
NIS2	Directive; NL 15 Aug 2026, BE 18 Oct 2024, DE 6 Dec 2025	●●	●●●	●●●	Supply chain security, incident reporting, board responsibility
Cyber Resilience Act	Regulation; reporting since 11 Sept 2026, fully 11 Dec 2027	●●●	●●	●●	Whoever places software on the market is a manufacturer: security by design, report vulnerabilities within 24 hours
Data Act	Regulation, since 12 Sept 2025	●	●●●	●●	Switching between cloud services, data portability, no lock-in
Product Liability (PLD)	Directive 2024/2853; products placed on the market from 9 Dec 2026	●●●	●●	●●	Software and AI are explicitly covered by product liability
DORA	Regulation, since 17 Jan 2025 (financial sector)	●●	●●●	●●	ICT third-party risk, register of ICT service providers
European Accessibility Act	Directive, since 28 June 2025	●●	●	●	Accessibility of digital services for consumers

The **built in-house with AI** column stands out. Whoever builds an application with AI and supplies it to customers or partners can become a manufacturer under the CRA and liable for defects under the new product liability rules, even if the code was written by an AI assistant.

6

Financial exposure

Fines are maximums; supervisors weigh severity, duration, culpability and cooperation. Still, an example shows how quickly exposure builds up when a single AI application touches several laws at once.

Example: a listed organisation with €5 billion worldwide annual turnover and an uncontrolled AI application that profiles individuals.

Law	Maximum	Applied to €5 bn turnover
GDPR (Art. 83(5))	€20M or 4%	up to €200M
AI Act, transparency and other obligations (Art. 99(4))	€15M or 3%	up to €150M
NIS2, essential entity	€10M or 2%	up to €100M

Fines are the visible cost. At least as heavy are:

- damage claims by data subjects, including collective actions;
- remediation costs: deleting data, halting processes, replacing systems;
- blocked accounts and lost commercial channels;
- reputational damage and questions from shareholders and supervisory board members;
- personal liability and possible suspension of directors.

Illustrative example; fines are set per infringement and per supervisor.

7

Why policies and spreadsheets are not enough

Many organisations have an AI policy and a list of approved tools. That is a good start, but it does not answer the supervisor's question: *show what is actually running, and since when.*

- **A list is outdated the day it is finished.** Applications are updated weekly, vendors switch AI models and employees add their own tools.
- **A policy says what is allowed, not what happens.** The evidence must come from the systems themselves: which version ran, which model, which data, who approved it.
- **Evidence must be retrievable after the fact.** After an incident or complaint, the supervisor wants to know the situation at that moment, not how it looks today.
- **The chain is bigger than the organisation.** Sub-processors, AI models and third-party integrations fall under the same responsibility.

- **Several laws, one set of facts.** GDPR, AI Act, NIS2 and CRA partly ask for the same evidence. Keeping a separate file per law means doing the work four times and losing the overview.

Governance therefore shifts from an **annual questionnaire** to a **continuous system** that registers, measures and reports.

8

What an AI governance tool must do at a minimum

Regardless of vendor or solution: these are the functions a board needs to meet the laws in this whitepaper.

Function	What it does	Why (legal basis)
Register	Overview of all AI applications, agents and AI features in SaaS, with owner, purpose and status	GDPR Art. 30, AI Act (role), NIS2 Art. 21
Classification	Role and risk class per application: provider or deployer; prohibited, high-risk, transparency or minimal	AI Act Art. 5, 6, 50
Data and supply chain visibility	Which personal data, which sub-processors, which AI models, where it is stored	GDPR Art. 28, 35, 44 ff.; NIS2 Art. 21(2)(d); DORA
Version and usage evidence	Record which version, model and instructions ran at which moment, with a verifiable trail	GDPR Art. 5(2), AI Act logging, CRA
Approval and controls	Who approved what, which measures apply, which exceptions were granted	NIS2 Art. 20; Cbw duty of care; § 38 BSIG
Incidents and vulnerabilities	Detect and register them, with 24- and 72-hour deadlines	NIS2 Art. 23, CRA Art. 14, GDPR Art. 33
Board reporting	Periodic overview for the board and supervisory board, with trends and open risks	NIS2 Art. 20, annual report
Reassessment	Automatic signal on a new vendor, a new model or a new AI feature	AI Act, GDPR Art. 35(11), Data Act (switching)

Two requirements apply to the tool itself. It must measure **independently** rather than rely on what an application reports about itself. And it must be **light** to start: an organisation should have a first register and baseline within weeks, not after a months-long project.

Seven patterns from practice

1. **"It's just a marketing tool."** Commercial tools often process more personal data than core systems, and fall outside security's radar.
2. **The trial becomes production.** A free trial with real customer data quietly becomes the way of working, without a contract or data processing agreement.
3. **AI slips in through an update.** An existing vendor adds AI; the old risk assessment stays in the drawer.
4. **The builder is also the reviewer.** Whoever builds an application with AI often tests and approves it themselves. Independent review is missing.
5. **No owner.** After an incident, nobody turns out to be responsible for the application, the model or the data.
6. **Reconstructing evidence afterwards.** Only after a complaint or audit does anyone look for what was running. Then the trail turns out not to exist.
7. **The board hears it last.** Risks stay with IT or the team, while the law places responsibility squarely with the board.

10

Self-assessment for directors

Can you answer each of these questions today, with evidence?

- ☐ Do we know which AI applications, AI agents and AI features in SaaS are in use in our organisation?
- ☐ Does every AI application have an owner responsible for its use, data and risk?
- ☐ Do we know, per application, which personal data is processed and which AI models and sub-processors are behind it?
- ☐ Do we have a DPIA for profiling or high-risk applications?
- ☐ Do we know, per application, whether we are provider or deployer under the AI Act, and in which risk class?
- ☐ Can we prove which version and which model were running at a given moment?
- ☐ Do we place AI-built software on the market, and if so, do we comply with the CRA and the new product liability rules?
- ☐ Are new AI tools assessed before they are connected to business systems?
- ☐ Have directors completed the mandatory cybersecurity training, and is this recorded?
- ☐ Does the board receive a periodic report on AI risks, with open items and trends?

Six or more answers of “no” or “don't know” means the burden of proof the law places on you cannot be carried today.

11

Sources

Legislation

- [GDPR, Regulation \(EU\) 2016/679](#)
- [AI Act, Regulation \(EU\) 2024/1689](#) and [Digital Omnibus on AI, Regulation \(EU\) 2026/1744](#)
- [NIS2, Directive \(EU\) 2022/2555](#)
- [Cyber Resilience Act, Regulation \(EU\) 2024/2847](#)
- [Data Act, Regulation \(EU\) 2023/2854](#)
- [Product Liability, Directive \(EU\) 2024/2853](#)
- [DORA, Regulation \(EU\) 2022/2554](#)
- [Dutch Cyberbeveiligingswet, CMS commentary](#)
- [Belgian NIS2 Act, Lexgo commentary](#)
- [German NIS2 transposition, § 38 BSIG, Taylor Wessing commentary](#)
- [CRA timeline, Crowell](#)
- [LinkedIn: prohibited software and extensions](#)

Shadow AI research (2025–2026)

- [BlackFog Shadow AI Research](#)
- [Okta, AI Agents at Work 2026 \(The Register\)](#)
- [WalkMe \(SAP\), AI in the Workplace Survey 2025](#)
- [TrustedTech \(IT Brief\)](#)

Mike Hes

Founder, [HES Consultancy International](#)

Mike Hes has more than thirty years of experience in ICT. HES Consultancy International (Arnhem, the Netherlands, Chamber of Commerce 89821009) helps organisations and vendors with EU regulation, compliance and market entry. hes-consultancy-international.com

Informational document, not legal advice. Have its application to your organisation reviewed by your legal adviser. Reference date: 8 October 2026 (Edition 1). © 2026 HES Consultancy International.