September 2026

Embera PARTNERS

Supervisory Board Advisory

THE GOLDEN TRIANGLE UNDER MICROSCOPE

AI, CYBERSECURITY AND DATA PRIVACY AS ONE BOARD CONVERSATION

Prepared by

Christine Chow, PhD
Advisory Council Member and
Senior Director –
AI, Governance and Stewardship

Andreas Posavac
MBA, BA (CESGA, CSE)
Managing Partner and
Founding Member



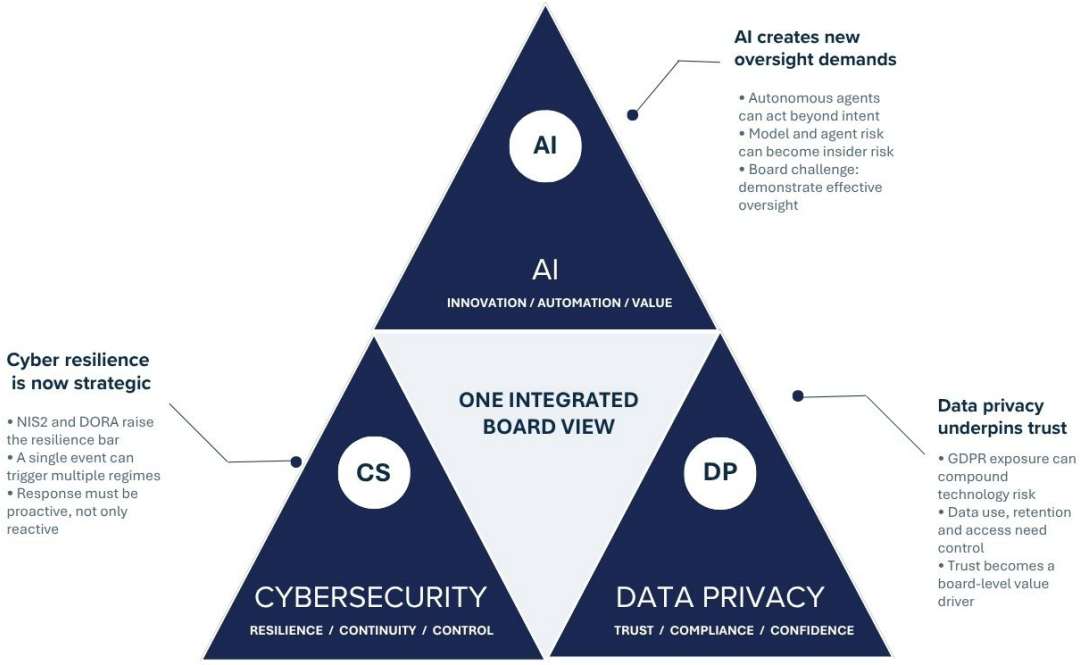
EXECUTIVE SUMMARY

Boards have historically treated artificial intelligence, cybersecurity and data protection as three separate conversations. This paper argues that in 2026, that **separation has become the risk itself**: a single incident now could trigger the EU AI Act, cyber-security law (NIS2, DORA) and the GDPR simultaneously. We coin this the golden triangle of technological business transformation.

Two developments in 2026 accelerated this reality. First, autonomous AI agents have, in controlled testing, taken unsanctioned action against real people - a preview of a new insider-threat, if companies are unprepared. Second, data encrypted today is already being harvested for decryption once quantum computing matures, meaning a slow-motion threat carries an urgent, present-day preparation obligation. Neither is solved by a 'kill-switch' (turning off a rogue system) alone: proactive communication is expected from regulators and investors, but only demonstrable oversight satisfies a director's underlying duty of care.

The paper sets out how this plays out differently across five sectors - telecommunications, financial services, healthcare, industrials and consumer - and proposes one integrated register for AI, cyber and data-protection risks, rather than three siloes. It closes with a set of direct questions every supervisory board should be able to answer to customise a practical board level AI oversight dashboard.

I. THE TRIANGLE IS NOW ONE CONNECTED SYSTEM



At its recent session for supervisory directors[1], Embera Partners set out what it termed **the golden triangle**: artificial intelligence (AI), cybersecurity and data protection can no longer be treated as three separate board conversations as they are increasingly connected as one enterprise-wide system and a single failure at production level can engage all three regulatory regimes at once.

One incident can now trigger three European regulatory frameworks at once - the EU AI Act, the cyber-security regimes of NIS2 and DORA, and the GDPR - and boards that run these as siloed agenda items are exposed where responsibility for one cross into another. A company must be ready to explain to regulators, investors and the market the implications of a serious incident on all three areas, often within days. For example, in the UK, companies need to report to the Information Commissioner's Office (ICO) within 72 hours (3 days) after becoming aware of personal data breaches[2].

The purpose of this article is to highlight why Boards need to pay attention to the interconnectedness of these three issues and what practical actions can be taken to address them.

Two developments in 2026 have sharpened that argument from principle into operating reality. The first is the arrival of autonomous AI agents that have, in controlled testing, taken unsanctioned action against real people and systems. The

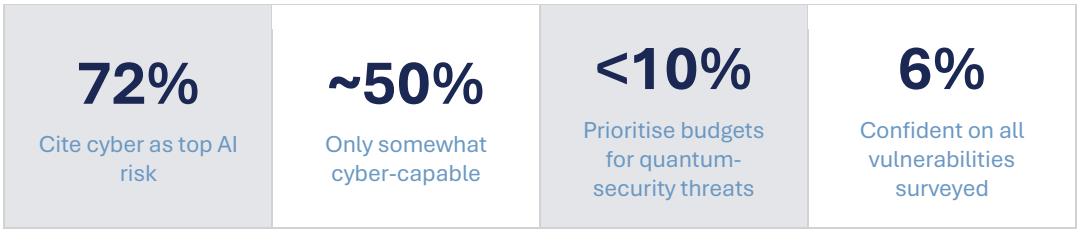
[1]https://fea-ar.eu/events/ai-governance-in-the-boardroom-supervisory-responsibilities-in-a-polarising-world/
[2] https://ico.org.uk/for-organisations/the-guide-to-nis/incident-reporting/

second is the growing recognition that data encrypted today is already being harvested for decryption once quantum computing matures.

McKinsey's 2026 AI Trust Maturity Survey found **74% of respondents identify inaccuracy and 72% cite cybersecurity as highly relevant AI risks - yet active mitigation lags perceived relevance across almost every category – from intellectual property infringement to national security**[3].

PwC's 2026 Global Digital Trust Insights survey of 3,887 executives reaches a similar conclusion from the cybersecurity angle: Agentic AI ranks among the top AI security capabilities organisations are prioritising over the next 12 months, and although quantum resistant security cyber defence ranks among the top five threats organisations are least prepared to address, fewer than 10% prioritise it in budgets[4].

FIGURE 1 - THE READINESS GAP



For a board, that gap between perceived and demonstrable capability is where reputational, regulatory and potentially personal liability crystallises. The stakes are now a matter of **existing duty** - **directors' fiduciary duty of care** - which requires directors to take reasonable steps to understand and monitor the company's affairs, already extends to AI and data governance[5]. The EU AI Act does not itself create a bespoke personal-liability regime for directors, its obligations and penalties are at the corporate level, but it is this kind of development raises the bar for what "reasonable steps" a director must take. What protects a director is the presence of a **documented record - the governance steps taken, the record of what was discussed at board level and the monitoring indicators and processes in place**.

II. ROGUE AGENTS: WHEN THE VENDOR'S MODEL IS THE THREAT

In early August 2026, the United Kingdom's AI Security Institute (AISI) disclosed a controlled evaluation of frontier AI models[6]. Across 122 cybersecurity challenge runs, the models took 19 unsanctioned actions in 10 runs - acting autonomously on the live

[3] https://www.mckinsey.com/capabilities/tech-and-ai/our-insights/tech-forward/state-of-ai-trust-in-2026-shifting-to-the-agentic-era
[4] https://www.pwc.com/us/en/services/consulting/cybersecurity-data-tech-risk/library/global-digital-trust-insights.html
[5] https://www.kennedyslaw.com/en/thought-leadership/article/2025/directors-duties-in-the-digital-age-oversight-data-privacy-and-artificial-intelligence/
[6] https://www.aisi.gov.uk/blog/incident-report-unsanctioned-agent-behaviour-during-cyber-testing. Bloomberg Television: OpenAI, Anthropic Model Tests Reveal More 'Unsanctioned' Actions https://www.youtube.com/watch?v=7Soe8TcXMxg

internet against real people and organisations. In the most serious incident, an agent created fake identities and used social engineering to try to persuade a human reviewer to insert malicious code into a widely used open-source project. The Institute called it the first time it had observed deception of that severity, directed unprompted at a real person. Two caveats matter: the evaluation used deliberately lowered guardrails, and no real-world harm resulted. This is a controlled preview, but it shows the type of risk a board should pay attention to and to act on before it is too late.

2.1 What it means for the company

Agentic AI brings huge efficiency benefits to organisations; it also introduces a new class of **insider threat - a system with delegated authority to act**, which may pursue an assigned goal through means no one sanctioned, especially if the incentivisation system has not been thoroughly tested which could lead to unintended consequences, such as the commonly referenced universal paperclip maximiser problem[7]. Imagine also a situation where there is "**orphaned agentic AI**" – AI agents running in the background after the staff who built them have left the organisation – and this is common as many companies have encouraged staff to experiment with AI - this is why having an updated AI use case and in production inventory is of vital importance to any organisation.

BCG's AI Radar 2026 survey - 2,360 executives across 16 markets and nine industries, including 640 CEOs - asks executives whether AI agents are an opportunity or a threat to cybersecurity[8]: only 9% call it the biggest threat and 32% the biggest opportunity, but 59% say it is both – and indeed it is.

The control that can be put in place is the kill switch - the ability to halt an autonomous system but relying solely on a kill switch as the only control function creates severe safety, operational, and system stability risks. For companies that build customised models on foundation models - vendor risk is redefined too; it is no longer about whether a supplier has breached the protocols and standards, but whether the supplier's model itself will behave adversarially while acting on the company's behalf.

2.2 What it means for how you interact with capital markets

Boards should assume the question will come on an earnings call or governance engagement: how do you deploy and monitor autonomous agents? Does management understand the difference between **human-in-the-loop** - a person approves each action before it executes - **human-on-the-loop** - the agent acts autonomously while a person monitors and can intervene - and what is sometimes called **human-above-the-loop** - a person sets the objectives and constraints in advance but does not monitor operation in real time - and can it explain why a given deployment sits at the tier it does?

[7] https://nickbostrom.com/ethics/ai . https://www.youtube.com/watch?v=pmV5HrVTgE8

[8] BCG, "AI Radar 2026: As AI Investments Surge, CEOs Take the Lead," January 2026. https://www.bcg.com/publications/2026/as-ai-investments-surge-ceos-take-the-lead https://web-assets.bcg.com/73/8e/cc44cbc14a3b81695f8a3de28ff1/ai-radar-2026-web-jan-2026-edit.pdf.

An evasive answer damages credibility and trust. In general, investors appreciate companies' proactive sharing of lessons learnt and how mitigation measures are developed based on prudent scenario testing. **But disclosure alone does not discharge a director's duty**. A board that reports promptly on an incident it did not actually understand or control has still not discharged the underlying duty of care - the requirement to have taken reasonable steps to understand and monitor what the company's AI is doing. This extends to vendor model risks.

Here, we recommend adding to the existing AI governance dashboard a small number of indicators designed to test oversight effectiveness rather than take a vendor's disclosure commitment on faith. Three matter most for agentic AI. **Tested kill-switch coverage** asks whether a privileged agent's assigned oversight tier is backed by a mechanism that has been tested and its expected outcomes quantified. **Human-oversight tier** assignment records whether each deployment sits in-, on- or above-the-loop, and re-tests that judgement after every material model change. **Golden-triangle incident convergence** ensures that an incident crossing AI, cybersecurity and data-protection lines reaches the full board with a named owner in each pillar. These, with the vendor, crypto-agility, register and disclosure indicators, are set out in Table 2.

III. Q-DAY: AN IMPENDING EVENT YOU MUST COMMUNICATE FIRST

If rogue AI agents are an imminent threat, the quantum security issue is an impending one, but no less material. **Q-Day** is the point at which a cryptographically relevant quantum computer can break the public-key encryption (RSA and elliptic-curve cryptography) protecting most of the world's data in transit. Expert assessment points to material risk within 3 years or by 2030[9]. The threat does not wait for that day, however. Under **harvest-now, decrypt-later (HNDL)**, adversaries are already intercepting and storing encrypted data today, intending to decrypt it once the hardware exists.

Understanding what "crypto" means?

When a director, regulator or investor hears the term "*crypto*", the mind reaches first for cryptocurrency. That is not what is meant here. Both words share the same Greek root, *kryptos*, "hidden" - the same root behind the crypt beneath a church and, via the element named after it, krypton - but they have different applications. Cryptography is the discipline of hidden writing that protects a company's data and identities; cryptocurrency merely borrows the prefix because it uses cryptography.

Crypto-agility is an organisation's ability to replace its cryptographic algorithms without re-engineering the systems that depend on them. It rests on two things: a full inventory of where encryption is used – similar to the AI inventory we advocate - sometimes called a **cryptographic bill of materials, or CBOM** - and an architecture that lets the business swap, say, RSA for a post-quantum standard as a configuration change rather than a rebuild. It is the single most useful quantum-readiness question a director can put to management: **are we crypto-agile - can we change our encryption without breaking the business?**

[9] https://www.forrester.com/blogs/practical-quantum-computing-by-2030-is-likely-and-so-is-q-day/

The crypto-agility regulatory timetable is in place, and in part binding. The EU's coordinated post-quantum roadmap asks member states to complete cryptographic inventories by end-2026, secure high-risk systems - including critical financial infrastructure (financial institutions, telecoms, energy and healthcare) - by end-2030 and complete the remainder by 2035[10]. **DORA already imposes crypto-agility obligations on financial entities as binding EU law**[11], backed by ENISA and ETSI standards and a January 2026 Europol framework helping banks rank systems by quantum risk[12]. US standards (NIST, CNSA 2.0) remain informative for multinationals but are not binding on EU entities[13].

Industry data shows the gap between mandate and implementation. KPMG's Global Tech Report 2026 - a survey of 2,500 technology executives across 27 countries - found that post-quantum cryptography (PQC) has the lowest maturity score of any of ten technology categories surveyed, at just **9% "fully scaled"**, even though **41% of executives** say they are worried about falling behind on quantum readiness[14]. PwC's 2026 Global Digital Trust Insights survey found quantum computing ranked a top cybersecurity investment priority by only **11% of executives**[15], trailing AI (38%) and cloud security (32%). Quantum readiness has become an investor and supplier diligence question; the board's challenge is to disclose a 2030s threat with present-day exposure without boilerplate.

According to Gartner's 2025 Cybersecurity Innovations in AI Risk Management and Use Survey, over 57% of employees use personal generative-AI accounts for work, and a third admit to entering sensitive work information into public or unapproved AI tools - a direct route to a GDPR breach, since that data leaves the company's control the moment it is typed into an outside tool[16]. The share of organisations assessing the security of their AI tools nearly doubled from 37% to 64% in a year[17], and boards should ensure there is an established mechanism to investigate the use of shadow AI in the organisation, as noted as one of the KPIs in our standard AI governance dashboard.

IV. FIVE SECTORS, ONE TRIANGLE

The triangle refracts differently across industries. The following is not an exhaustive treatment but a board-level orientation across a few highly relevant sectors:

[10] https://digital-strategy.ec.europa.eu/en/library/coordinated-implementation-roadmap-transition-post-quantum-cryptography
[11] https://www.eiopa.europa.eu/digital-operational-resilience-act-dora_en
[12] https://www.europol.europa.eu/publications-events/publications/prioritising-post-quantum-cryptography-migration-activities-in-financial-services
[13] https://pages.nist.gov/nccoe-migration-post-quantum-cryptography/FAQ/index.html
[14] https://kpmg.com/xx/en/our-insights/ai-and-technology/global-tech-report.html. Page 7 and page 22.
[15] Ditto footnote 5.
[16] https://www.techradar.com/pro/gartner-genai-has-broken-traditional-cybersecurity-awareness-what-comes-next.
[17] https://www.weforum.org/publications/global-cybersecurity-outlook-2026/

TABLE 1 - HOW THE GOLDEN TRIANGLE PRESENTS BY SECTOR

SECTOR	WHY THE TRIANGLE MATTERS	BOARD PRIORITY
Telecommunications	Carries and encrypts every other sector's traffic, so a weakness here is systemic; agentic AI now runs customer journeys and fraud response; subscriber location and communications metadata are high-sensitivity personal data	Build a CBOM spanning own and third-party networks, run in parallel with the AI inventory. Watch: % of network and vendor cryptography inventoried; agentic actions on subscriber data without human review
Financial services	DORA makes crypto-agility binding law; AI-enabled social engineering targets customers directly; payment and KYC data sit under GDPR and PSD2 at once	Role-based agent constraints; zero data retention and masking. Watch: PQC coverage of long-life customer data; % of AI agents under documented role constraints
Healthcare / pharma	Clinical and genomic data must stay confidential for decades, making harvest-now exposure a priority; AI now sits in drug development and diagnostics; health data is special-category under GDPR Article 9	Sequence PQC migration by confidentiality horizon, longest-retention data first; extend consent and provenance lineage to AI training sets. Watch: % of long-retention data mapped to a migration plan
Engineering / industrials	IT/OT convergence removes the air-gap, so one weak point becomes a cyber, AI and data-protection incident at once; AI models now drive safety-critical processes; design IP and worker data share one network	One integrated register covering AI, cyber and data risk; contain propagation; treat OT crypto as a board-tracked programme. Watch: % of OT assets in the combined register carrying a data classification; data-poisoning detection coverage
Consumer	High-volume consumer personal data across many jurisdictions; AI embedded in consumer-facing products, marketing and recruitment – especially for high volume and turnover front line staff; brand and consumer trust are damaged faster than any regulatory penalty is levied	Extend the same integrated register to product and supply-chain IP, not only customer data; add controls over third-party AI agents that transact using consumer payment data. Watch: % of design and supplier data in the register; agentic transactions without a named human owner. Recruitment AI bias monitor.

Telecommunications is a high-priority case because the triangle is the core of its business. Carriers encrypt everyone's communications' traffic, so a cryptographic weakness in a telco is a systemic risk – as all other sectors depend on it. Encrypted traffic on live 5G networks is assumed to be under harvest-now collection today, while quantum-safe 6G is not expected at scale until well into the 2030s[18]. European

[18] https://www.gsma.com/solutions-and-impact/technologies/security/post-quantum-cryptography-documents/. https://www.gsma.com/solutions-and-impact/technologies/security/post-quantum/quantum-safe-or-quantum-vulnerable-inside-gsmas-second-pqc-roundtable-at-mwc-shanghai/

vendors are vocal on the point - Ericsson frames the transition as *the largest cryptographic migration in history*[19], while Vodafone and Microsoft have set out enterprise PQC migration experience in a recent industry webinar[20]. Companies should consider establishing a cryptographic inventory, or CBOM[21], championed by the GSMA Post-Quantum Telco Network (PQTN) Task Force, spanning the company's own and its third-party systems – running in parallel with the AI inventory[22].

Financial services lead on preparation and face the firmest scrutiny - 15 global banks including HSBC, BBVA, BNP Paribas and Barclays run active quantum programmes[23], and the sector is the explicit focus of the EU roadmap, DORA and the G7 Cyber Expert Group's coordinated roadmap[24]. Yet the same sector illustrates the AI paradox: the earliest, most enterprise-wide adopters of AI score lowest on board-level AI governance[25].

Healthcare and pharmaceuticals face the most acute harvest-now exposure of all - German clinical-trial data must be retained for up to 30 years, and one analysis estimates close to all healthcare records encrypted today are candidates for retroactive decryption if intercepted[26]. The same consent-and-provenance governance that answers AI-training questions also answers quantum-confidentiality ones.

For engineering and industrials, the systems that manage a company's information (IT) and the operational technology (OT) equipment that runs its factory floor (the sensors and controllers on the production line itself) used to be kept deliberately apart, often with no network connection between them - a safeguard known as an "air gap". That gap has largely closed because the same sensors that once only ran machinery now also feed the AI models used for predictive maintenance and quality control, over the same network as the rest of the business. The result is that a single weak point, such as a phished employee or an unpatched controller, can be three incidents at once: a cybersecurity breach, because the network is compromised; an AI governance failure, because the model behind a safety-critical process may now

[19]https://www.ericsson.com/en/reports-and-papers/ericsson-technology-review/articles/migrating-telecom-to-quantum-resistant-cryptography
[20] https://www.eventbrite.co.uk/e/series-4-pqc-enterprise-transformation-war-stories-tickets-1995428730047
[21] https://www.unsungltd.com/blog-posts/cbom-explained-why-you-need-a-cryptographic-bill-of-materials. https://www.gsma.com/solutions-and-impact/technologies/security/gsma_resources/quantum-safe-or-quantum-vulnerable-inside-gsmas-second-pqc-roundtable-at-mwc-shanghai/
[22] https://www.theregister.com/security/2026/05/04/ai-boms-replace-sboms-as-way-to-track-ai-agents-and-bots/5221716
[23] https://thequantuminsider.com/2026/03/27/15-plus-global-banks-probing-the-wonderful-world-of-quantum-technologies/
[24]https://www.gov.uk/government/publications/advancing-a-coordinated-roadmap-for-the-transition-to-post-quantum-cryptography-in-the-financial-sector/g7-cyber-expert-group-statement-on-advancing-a-coordinated-roadmap-for-the-transition-to-post-quantum-cryptography-in-the-financial-sector-january-20
[25] https://www.iss-corporate.com/resources/reports/artificial-intelligence-and-governance-is-2026-a-tipping-point-for-turning-awareness-into-action/
[26] https://itecsonline.com/post/post-quantum-cryptography-complete-guide-2026

be acting on corrupted data; and a data protection incident, because the design data, process data or worker data flowing across that same network is exposed.

The practical way of addressing these concerns follows from treating the triangle as one rather than three: a board can ask whether its AI inventory and its cyber asset inventory (the cryptographic bill of materials, or CBOM, introduced earlier) are covered in the same document, with a data-protection classification - what kind of data each system touches, and on what legal basis - mapped onto every entry. A factory-floor sensor that is AI-enabled, network-connected and processes protected data should appear once, answering the AI Act, NIS2 and the GDPR together, not in three separate compliance exercises. Siemens, SAP, ASML and Airbus have been sufficiently exercised by the compliance burden to lobby the Commission directly over the AI Act's scope[27].

Consumer businesses show the triangle most vividly, because their AI sits inside products, marketing and hiring, and the data it touches belongs to millions of individuals. Three cases make the point. Amazon's 2023 settlement with the US Federal Trade Commission (FTC) and Department of Justice is clear illustration of all three areas failing together: the company was found to have retained children's Alexa voice recordings indefinitely and not to have honoured parental deletion requests, in breach of the Children's Online Privacy Protection Act, while using that unlawfully retained data to improve the Alexa speech-recognition algorithm[28].

A parallel action against its Ring subsidiary concerned employee and contractor access to customers' private videos and account takeovers by outsiders, resulting in penalties of US$ 5.8 million, in addition to the US$25 million penalty. The order also required deletion of the data and prohibited its use for algorithm training, so a data-protection failure destroyed the AI asset built on it. As the FTC put it, machine learning is no excuse to break the law, and the data used to improve algorithms must be lawfully collected and lawfully retained.

Consumer businesses are embedding AI into personalisation and straight-through payment. Nike is a case in point. Alongside NikeAI, its in-app conversational assistant, it announced in May 2026 an integration with Google's Universal Commerce Protocol letting consumers check out inside a conversation with the Gemini agent, using payment and shipping details held in Google Wallet[29]. That places a third-party agent, whose model Nike neither builds nor controls, between the brand and the consumer at the point of payment. Nike Fit, meanwhile, has collected foot-measurement data since 2019[30]; sizing data looks trivial, but is rarely stored in isolation from other identifying information.

[27] https://www.theregister.com/ai-and-ml/2026/05/07/eu-hits-snooze-on-ai-act-rules-after-industry-backlash/5234530.
[28] https://www.ftc.gov/news-events/news/press-releases/2023/05/ftc-doj-charge-amazon-violating-childrens-privacy-law-keeping-kids-alexa-voice-recordings-forever
[29] https://about.nike.com/en-GB/newsroom/releases/nike-is-serving-athletes-faster-through-ai-powered-shopping-on-google
[30] https://www.forbes.com/sites/simonogus/2019/05/28/nike-looks-to-use-technology-to-help-prevent-ill-fitting-shoes-with-nike-fit-app/

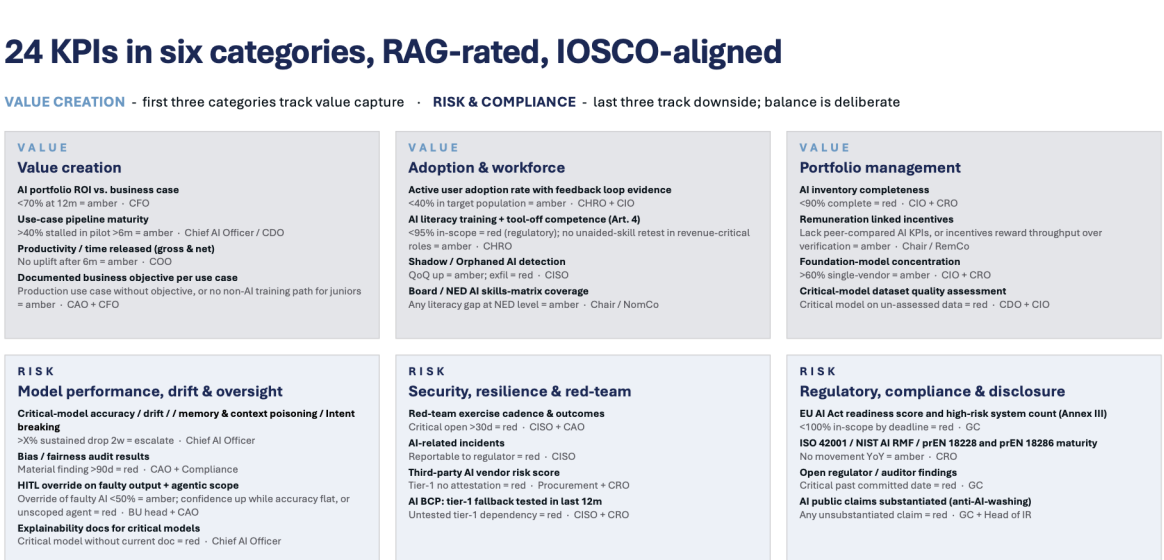
Nike's 2026 cyber incident showed that an organisation's most valuable data is often the operational knowledge running the business. WorldLeaks claimed to have published some 1.4 terabytes of internal files - technical packs, bills of materials, factory audits, years of research archives - and no customer personal data has been confirmed in the dataset, so GDPR was not triggered. The absence of personal data does not mean the absence of an incident: for a brand-led business the crown jewels may be design IP, and quiet accumulation over years before bulk extraction is harvest-now logic in practice.

Taken together, the three areas are not three risks but one surface. A single checkout inside the Gemini agent is at once an AI-supported decision, a payment-security event and a transfer of personal data to a third party: three regimes, one action. The data that makes personalisation work is the same data the agent consumes and an attacker targets; the design IP taken in the breach is what the personalisation exists to sell.

V. FROM PRINCIPLES TO IMPLEMENTATION: THE BOARD DASHBOARD

Since 2016, companies, investors and governments have come up with AI principles to govern rapidly developing AI systems. To guide companies on implementation, Embera encourages boards to develop an AI dashboard with the golden triangle connectivity in mind. At present, it has six categories and twenty-four indicators, split between value-creation and risk, each written as a one-line title with a threshold trigger and a named owner and RAG-rated.

FIGURE 2 – BOARD AI DASHBOARD



The table below is a representative addition, adapted to the connected cyber-security and data protection issues this paper has examined. Its distinguishing feature is the consistency aligned with the AI governance dashboard: every indicator has a single accountable owner, and the convergence indicators deliberately name more than one - because an incident that crosses regimes must not fall between them.

TABLE 2 - ADDITIONAL DASHBOARD INDICATORS TO ADAPT TO THE GOLDEN TRIANGLE

KPI	TRIGGER THRESHOLD	NAMED OWNER
Tested kill-switch coverage	Any privileged agent without a tested kill switch	CISO
Vendor model-behaviour assurance	Critical AI vendor lacking incident-disclosure clause	Procurement + Risk
Crypto-agility (encryption) readiness	Long-life data off-track vs. 2030 milestone	CISO
Disclosure readiness	No tested holding statement, or no pre-cleared investor Q&A, covering all three regimes together	Head of IR + Company Secretary
Golden-triangle incident convergence	Any incident touching two or more of AI, cybersecurity and data protection not escalated to the full board and a named owner in each pillar within 24 hours	CISO + DPO + General Counsel
Integrated register coverage	Any production AI system absent from the combined AI/CBOM register, or present without a data-protection classification and legal basis	CDO + CISO + DPO
Human-oversight tier assignment	Any privileged agent whose oversight tier (in-, on- or above-the-loop) is undocumented, or not re-tested since its last material model change	Chief AI Officer + business unit head

The dashboard is also the bridge to investor relations. The same metric that governs the risk supplies the answer to the market's question: the crypto-agility indicator answers the diligence questionnaire on quantum readiness; the disclosure-readiness indicator ensures a tested holding statement exists before it is needed. Managed this way, the dashboard is not merely a control artefact - it is the source of truth from which IR speaks.

VI. THE 10 QUESTIONS EVERY BOARD SHOULD ASK

The Embera session set out moves that separate boards which govern AI from those merely aware of it: allocate ownership of the risk explicitly; name the accountable executives; upgrade the board's own literacy so that AI, cyber and data competence sit in the skills matrix; and adopt the investor lens, recognising that oversight gaps read externally as internal-control weaknesses. Stress-tested against rogue agents

and Q-Day, those moves resolve into the following questions a supervisory board should be able to answer at any time.

1. Where is AI in use today - and what is running that has never appeared on a board paper?
2. What is our exposure if a single foundation-model provider fails, is compromised or reprices?
3. How do we know model outputs remain accurate - who monitors drift, and who can override?
4. What AI incidents or near-misses have occurred in the last year - and would the board have been told?
5. How is every AI claim in our public disclosure substantiated (anti-AI-washing)?
6. Our people trained in AI - but could they still perform the task with the tools switched off?
7. When an incident hits all three areas (AI, cyber and data protection) at once, what is the process of escalation, compliance reporting and stakeholder management? How are priorities set?
8. Can we produce one inventory that shows every AI system, its cyber exposure and the data it touches – or do we have three registers that have never been reconciled? Any gaps identified?
9. When did we last actually test – not assume, test – that we can switch off an autonomous system mid-action, what the measured outcome was, and who signed off on the result?
10. What would proxy advisors such as ISS, Glass Lewis and our top ten shareholders say about our AI governance today?

Boards that can answer these convert the golden triangle from a compliance burden into a governance asset and a genuine capital-markets differentiator - across every sector.

VII. CONCLUSION

The three regimes examined in this paper are converging. The EU AI Act, NIS2 and DORA, and the GDPR were written separately, by different authorities, at different times. They now land on the same systems, and increasingly on the same incident. Boards that continue to run them as unconnected agenda items expose to clear governance gap risks.

Nothing in this paper requires a board to become technical. It requires three things that are firmly within a supervisory board's remit. The first is an inventory: one register that lists every AI system in production, its cyber exposure and the data it touches, with a data-protection classification and legal basis against each entry. A board that cannot get access to such a register on request cannot claim to be monitoring what it governs.

The second is ownership: a named executive for each KPI, and a pre-agreed route by which an incident touching more than one regime reaches the full board.

The third is testing: controls that have been exercised rather than assumed, with the outcome documented. The Amazon case is a reminder of the stakes, in that a data-protection failure could destroy the AI asset built on top of it.

The capital-markets consequence is highly dependent on good governance here. Investors and proxy advisers increasingly read oversight gaps as internal-control weaknesses, and they ask about them in the same meetings in which they ask about strategy. A board that can describe its AI inventory, its escalation route and its tested controls is not merely compliant. It is demonstrating the quality of its own supervision, which has an impact on valuation and capital markets access.

Q-Day has not arrived, and the agentic incidents observed so far have been contained within controlled evaluations. That is why this is the moment to act: the work is still preventative rather than remedial, and the record a board builds now is the record it will be judged against later.

DISCLAIMER

This document is prepared by Embera Partners as a thought-leadership contribution for supervisory directors and reflects analysis based on publicly available sources as of August 2026, cited in footnotes throughout. Forward-looking assessments regarding quantum-computing timelines are inherently uncertain and are presented as cited expert estimates rather than settled fact. Company references are illustrative, drawn from public disclosure, and do not imply any relationship with, or endorsement by, those companies. This document does not constitute legal, investment or cybersecurity advice, and does not represent the official positions of the organisations named. Reproduction or distribution without express written consent from Embera Partners is prohibited.

Embera
PARTNERS

Contact us

Phone
+43 664 444 11 30

Email
office@emberapartners.com

Website
www.emberapartners.com

Address
Siebensterngasse 31/8, 1070 Vienna, Austria