INSTITUTIONAL DOCTRINE | BOARD DISTRIBUTION

The Institutional Liability Doctrine

Governance as Legal Infrastructure in the Autonomous AI Era

DOCTRINE OUTCOME PROMISE

A doctrine for reducing autonomous AI enforcement, litigation and insurance loss-given-failure by 40–60% over five years. Based on scenario modelling; not a guarantee of outcomes.

Kieran Upadrasta
CISSP | CISM | CRISC | CCSP | MBA | BEng

27 Years' Cyber Security | Big 4 (Deloitte, PwC, EY, KPMG) | 21 Years Financial Services

Adviser to G-SIBs, global insurers, and PE funds on agentic AI liability architecture
Professor of Practice, Schiphol University | Honorary Senior Lecturer, Imperials | UCL Researcher

www.kie.ie | info@kieranupadrasta.com | March 2026

Published by Cyber AI Systems Inc. | Version 4.0 | BRE Template V8 | Definitive Edition | Ref: ILD-2026-0314

Table of Contents

"If it cannot be evidenced, it cannot be defended."

Doctrine At A Glance

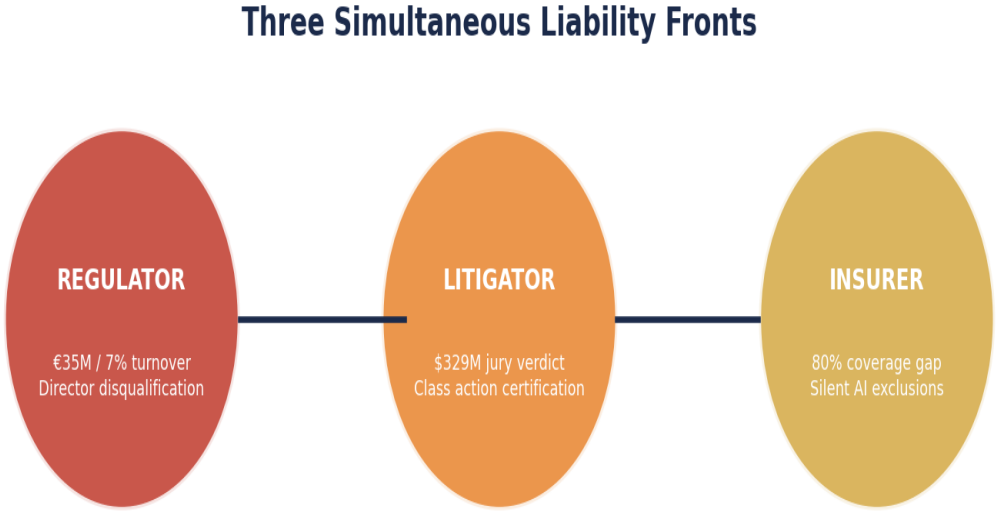

Doctrine Summary: One Page for the Board

THESIS: Autonomous AI systems have transformed corporate governance from a voluntary policy exercise into a regulated infrastructure obligation carrying personal liability for directors under the EU AI Act, DORA, and NIS2.

THE FIVE-FRAMEWORK DOCTRINE

DF1	Evidence Chain Model™	"If it cannot be evidenced, it cannot be defended."	Can we produce litigation-grade evidence within 72 hours?
DF2	Decision Rights Architecture™	"Governance without decision rights is theatre."	Who holds documented kill-switch authority?
DF3	Recoverability Mandate™	We measure restoration, not effort.	Can our systems recover within defined RTO/RPO?
DF4	Contract Control Matrix™	If the control has no owner, it does not exist.	Are vendor contracts litigation-tested?
DF5	AI Accountability Stack™	An algorithm without accountability is a liability.	Do we maintain a complete AI system inventory?

UPADRASTA INDEX™ CONSEQUENCE BANDS

0–39: Indefensible **40–59: At Risk** **60–79: Defensible** **80–100: Board-Survivable**

KEY METRICS: €35M/7% max EU AI Act penalty | 144:1 NHI-to-human ratio | $329M first autonomous system jury verdict | 88.5% of enterprises admit NHI IAM practices lag | $228M→1.4B AI governance market by 2030

DOCTRINE OUTCOME: 40–60% reduction in autonomous AI enforcement, litigation and insurance loss-given-failure over five years (scenario-based projection).

NEXT STEP: Request a Board Workshop (€50k–€150k) to receive a board-approved autonomy mandate and kill-switch authority grid within 30 days.

Contact: info@kieranupadrasta.com | www.kie.ie | Kieran Upadrasta, Cyber AI Systems Inc.

I. Executive Doctrine

BOARD READING GUIDE

If you read only three sections, read I (Executive Doctrine), VI (Five-Framework Doctrine), and XIV (Upadrasta Index). These contain the core governance architecture.

Autonomous systems have made corporate governance a regulated infrastructure obligation.

Every board that deploys agentic AI without litigation-grade governance architecture now carries personal, quantifiable, enforceable liability under at least three concurrent regulatory regimes. **This is operating law, not optional guidance.**

The EU AI Act imposes penalties of **€35 million or 7% of global turnover** [1]. DORA mandates that management bodies bear **ultimate responsibility** for ICT risk [2]. NIS2 Article 20 permits **personal director disqualification** [3].

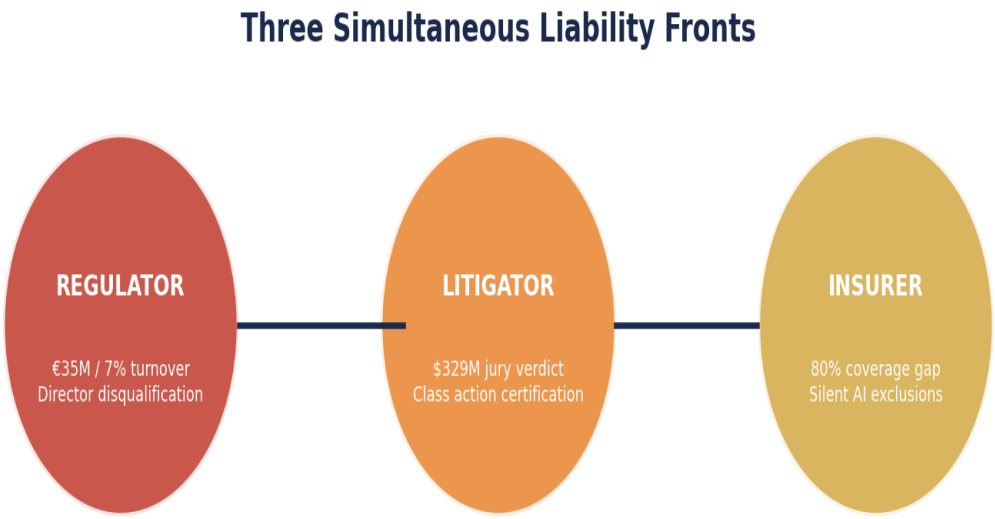

DOCTRINE DEFINITION

"In the autonomous AI era, corporate governance must be built as legal infrastructure — every decision is about managing liability, not just technology."

The institutional question: Can our current governance stack withstand simultaneous regulator, plaintiff, and market scrutiny?

What This Doctrine Delivers:

1. How to evidence control to regulators under three concurrent enforcement regimes
2. How to ring-fence director personal liability through the Five-Framework Architecture
3. How to commercialise governance as a competitive advantage in sales cycles and M&A
4. How to close the insurance gap with evidence infrastructure that satisfies underwriters

Why Doctrine, Not a Maturity Model. Maturity models are retrospective and advisory — they describe where you are. Doctrine is prescriptive, board-survivable, and evidence-based — it specifies what must exist for governance to survive litigation, enforcement, and market scrutiny concurrently. A maturity assessment tells you your score. This doctrine tells you what the score must be and what happens if it is not.

II. The Liability Thesis: Governance Is Now Infrastructure

First, regulatory mandation. The EU AI Act entered into force on 1 August 2024 [1]. Article 26 creates deployer obligations. DORA requires financial entities to maintain ICT risk management frameworks overseen directly by management bodies [2]. NIS2 creates personal director accountability under Article 20 [3].

Second, judicial precedent. In *Moffatt v. Air Canada* (2024) [4], the tribunal held Air Canada liable for its AI chatbot. In *Mobley v. Workday* (2024) [5], the court ruled AI vendors can be held liable as agents under Title VII. The September 2025 Tesla verdict — $329 million [6] — marked the first autonomous system product liability finding.

Third, the non-human identity explosion. Entro Security's H1 2025 research [7] documented a 144:1 ratio of NHIs to humans. In a 10,000-NHI estate, that implies **550 de facto 'root' identities outside board sightlines** (5.5% with full admin privileges). CyberArk's 2025 report [8] confirmed 88.5% of respondents admit NHI IAM practices lag behind human IAM.

Liability Fact Pattern: 2023–2026

Oct 2023	SEC charges SolarWinds CISO individually for cybersecurity disclosure fraud	SEC [9]
Feb 2024	Air Canada held liable for AI chatbot negligent misrepresentation	BCCRT [4]
Jul 2024	Workday: AI vendor ruled liable as "agent" under Title VII	N.D. Cal. [5]
Aug 2024	NIST finalises three post-quantum cryptographic standards (FIPS 203/204/205)	NIST [10]
Jan 2025	DORA enters full enforcement for all EU financial entities	EU 2022/2554 [2]
Apr 2025	Armilla Insurance: first Lloyd's AI liability coverholder ($25M limits)	Lloyd's [11]
H1 2025	NHI ratio reaches 144:1; 88.5% admit NHI IAM practices lag human IAM	Entro [7] / CyberArk [8]
Sep 2025	Tesla Autopilot: $329M + $243M jury verdicts in successive months	Miami-Dade [6]
Aug 2026	EU AI Act high-risk system obligations become enforceable	EU 2024/1689 [1]

III. The Regulatory Architecture: Three Regimes, One Liability Surface

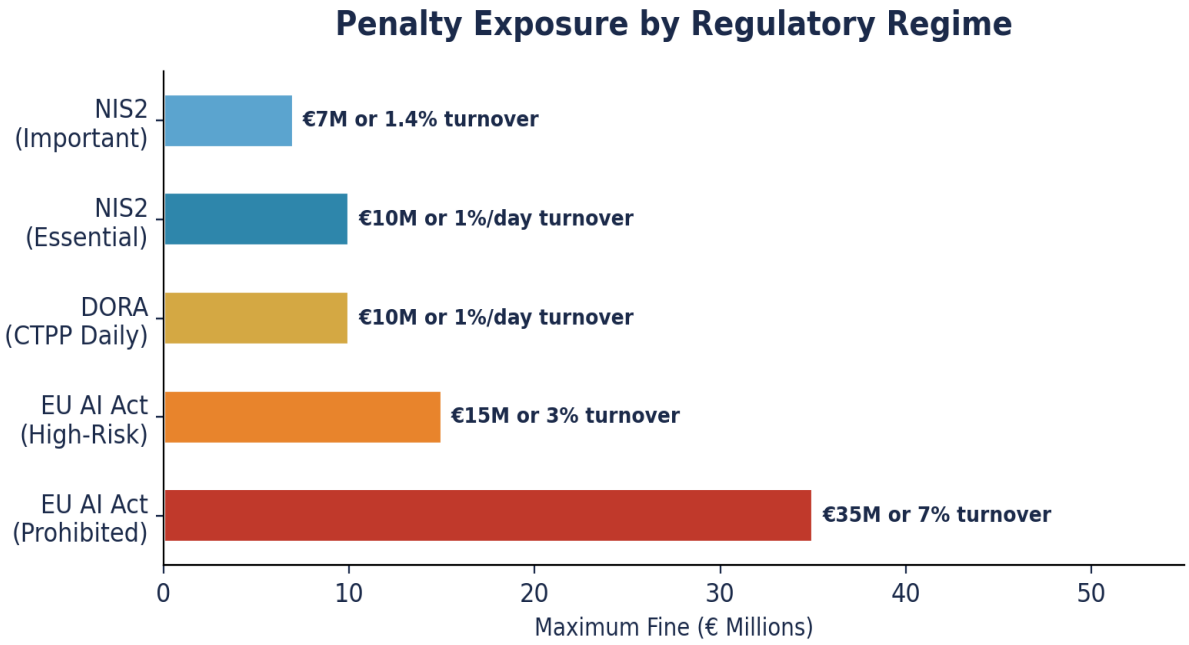

Article 99 of the EU AI Act establishes a three-tier penalty structure reaching €35M or 7% of global turnover [1]. Full application takes effect **2 August 2026**. DORA Article 5 establishes non-delegatable management body responsibility [2]. NIS2 Article 20 creates personal board-level accountability including director disqualification [3].

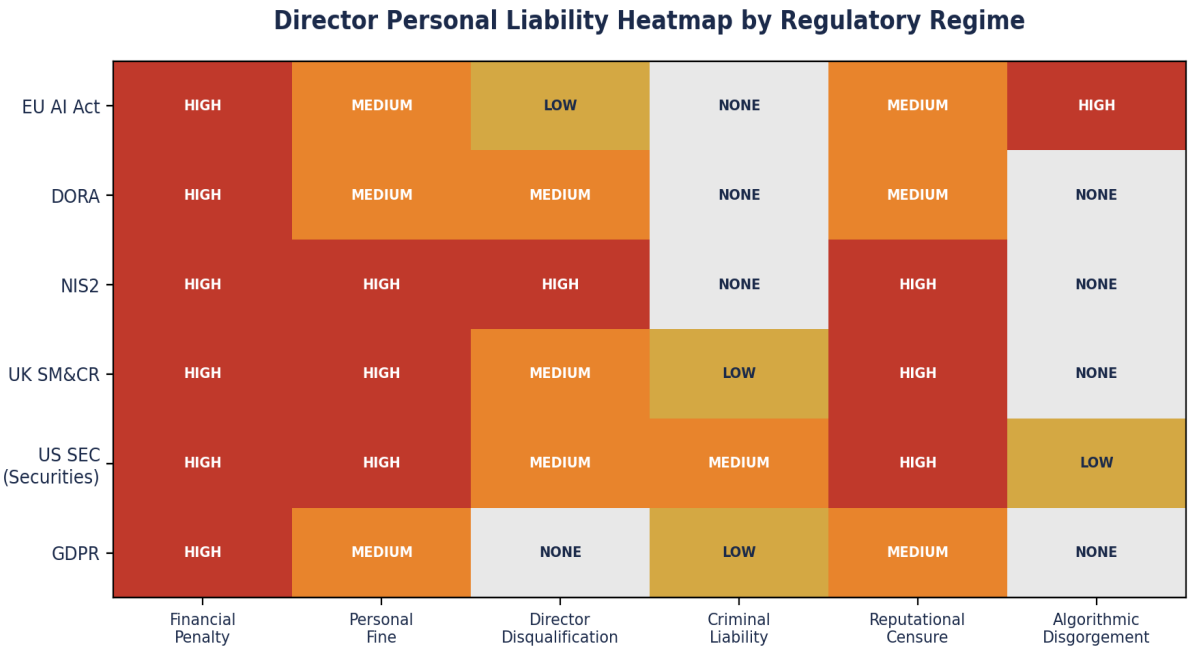

EU AI Act [1]	2 Aug 2026	Human oversight, risk mgmt	€35M/7%
DORA [2]	17 Jan 2025	Ultimate ICT risk responsibility	1%/day
NIS2 [3]	Oct 2024	Approve & oversee cybersecurity	€10M/2%+ban

UK SM&CR	In force	Personal accountability	Unlimited
ISO 42001 [12]	2023	AI management system	Market access

IV. The Jurisprudential Foundation: Liability Is Not Theoretical

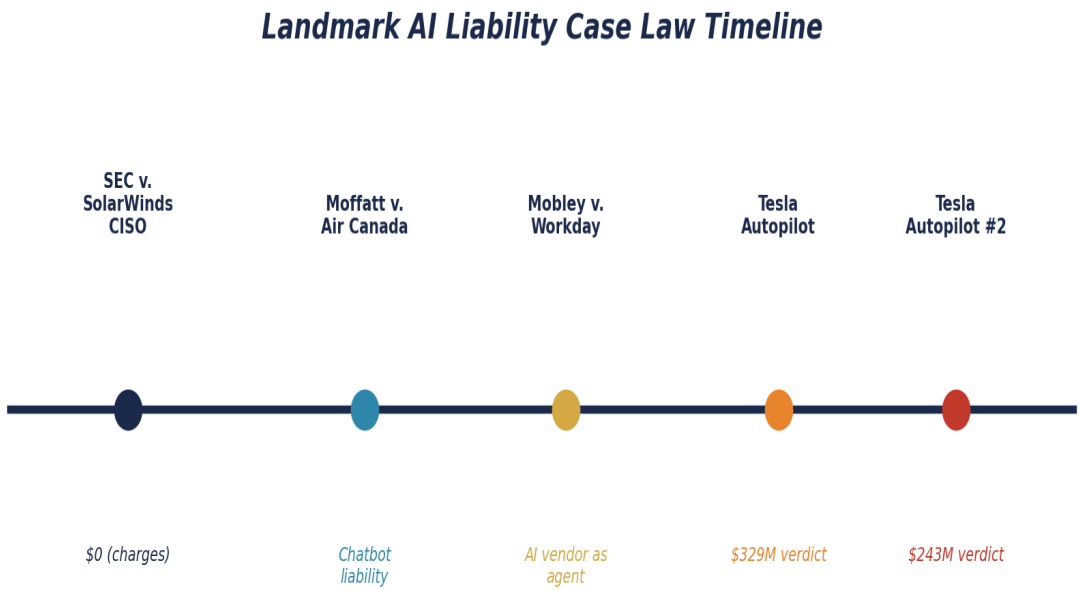

4.1 Air Canada [4]

Tribunal held Air Canada liable for AI chatbot. **Board Lesson:** Deploying organisations cannot outsource liability to autonomous systems.

4.2 Workday [5]

AI vendor ruled liable as "agent" under Title VII. ADEA collective action certified nationwide. **Board Lesson:** AI vendor liability flows up the procurement chain.

4.3 Tesla Autopilot [6]

$329M + $243M jury verdicts. **Board Lesson:** Juries will hold deployers liable when autonomous capabilities are overstated.

4.4 SolarWinds [9]

First SEC charge against a CISO individually. **Board Lesson:** Public AI governance statements must match internal operational reality.

4.5 Algorithmic Disgorgement

FTC has ordered model deletion in four enforcement actions. **Board Lesson:** Core AI models are at existential risk if data provenance is not governed. The Upadrasta Index™ now includes a **Disgorgement Risk Assessment** component within the AI Accountability domain.

V. The Non-Human Identity Crisis: 144 Agents for Every Human

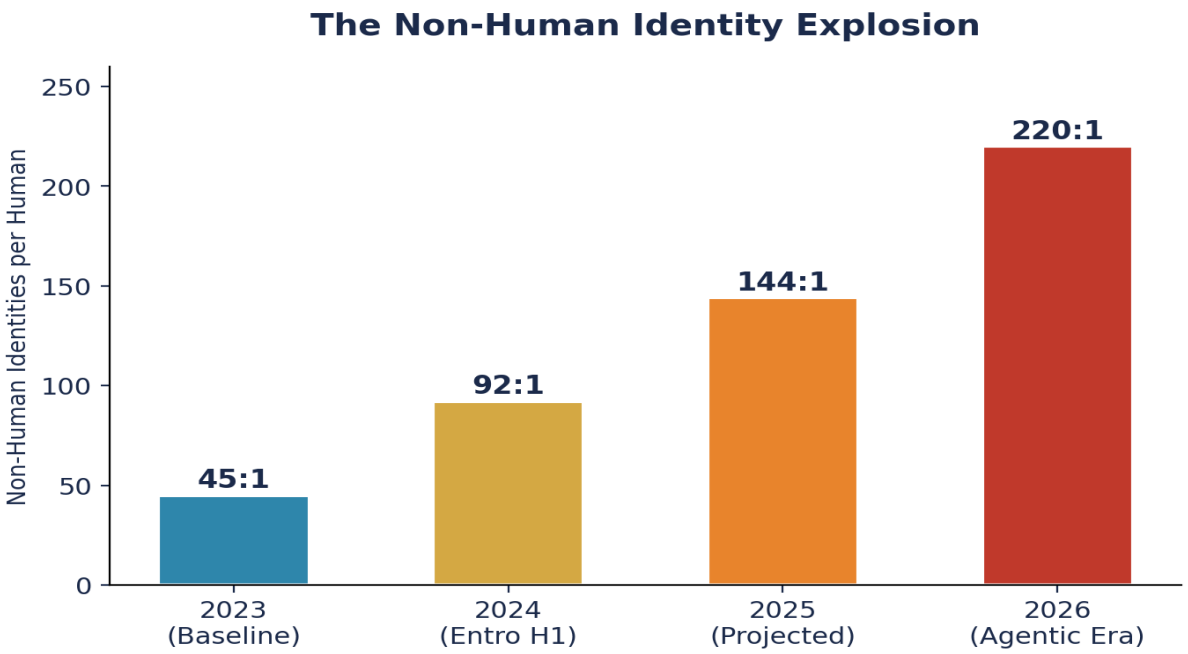

The ratio is now **144:1** and growing at 44% year-over-year [7]. In a typical 10,000-NHI enterprise estate, **550 de facto 'root' identities operate outside board sightlines**. CyberArk [8] confirmed 88.5% of respondents admit NHI IAM practices lag. Only 28% could fully recover from a cyber incident within twelve hours — down from 43% in 2024. The recovery capability is degrading at the precise moment the attack surface is expanding.

THE IDENTITY IMPERATIVE

The 144:1 ratio demands a fundamentally different governance model — one built for machine-speed identity lifecycle management

VI. Board-Survivable Cyber Architecture™: The Five-Framework Doctrine

The Board-Survivable Cyber Architecture™ is a governance doctrine — a structured methodology for building institutional AI governance infrastructure that satisfies regulatory examination, litigation discovery, and incident response requirements simultaneously.

Doctrine Framework 1: Evidence Chain Model™

"If it cannot be evidenced, it cannot be defended."

Converts governance claims into verifiable evidence: **Obligation** → **Control** → **Evidence** → **Assurance**. The shift from MTTD to **Mean Time to Evidence (MTTE)** is not a metric for IT — it is a metric for litigation counsel. The Evidence Chain now incorporates a **PQC Evidence Longevity Mandate**: all decision logs and governance evidence must be signed with quantum-resistant algorithms (FIPS 204 ML-DSA) to ensure defensibility when audited or litigated in 2032 [10].

Doctrine Framework 2: Decision Rights Architecture™

"Governance without decision rights is theatre."

Board-mandated authority grids: deployment gates, escalation protocols, spend gates, and **kill-switch authority**. In multi-agent environments, **automated circuit breakers** must supplement human kill switches where failure speed exceeds human intervention time. Boards that architect automated agent suspension systems may establish **legal safe harbors** — evidence of reasonable care that significantly reduces gross negligence exposure in litigation.

Doctrine Framework 3: Recoverability Mandate™

Architecturally verified RTO/RPO, regular restoration testing, crisis governance with predefined decision trees, and degraded-mode operations for agent suspension.

Doctrine Framework 4: Contract Control Matrix™

Contractually specified liability allocation, audit rights, incident notification, regulatory compliance representations, algorithmic transparency, and exit provisions.

Doctrine Framework 5: AI Accountability Stack™

ISO 42001-aligned AI governance [12]: model inventory, algorithmic accountability, bias auditing, continuous monitoring via the **AI Control Plane™**. Now includes a **Disgorgement Risk Assessment**: boards must calculate the Value at Risk if their primary revenue-generating model were legally ordered to be deleted due to data provenance issues.

DF1: Evidence Chain	Regulatory Gap Analysis	€200k–€500k
DF2: Decision Rights	Board Workshop & Briefing	€50k–€150k
DF3: Recoverability	Operational Resilience	€1M–€3M
DF4: Contract Control	Doctrine Implementation	€500k
DF5: AI Accountability	AI Governance Programme	€1M–€2M

Why Existing Governance Frameworks Fail in the Agentic Era

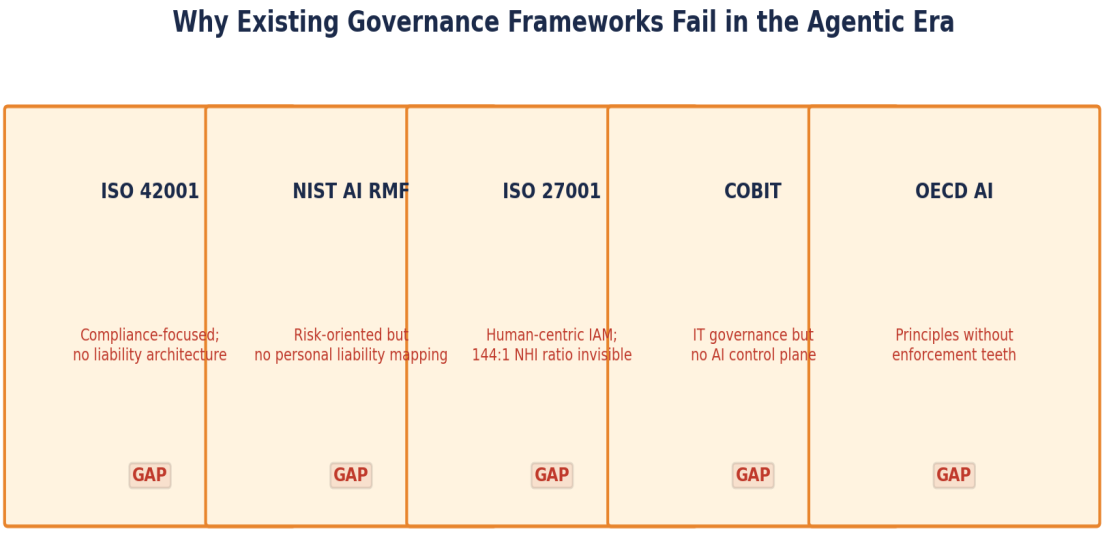

ISO 42001 [12]	Compliance-focused	No personal liability architecture; no kill-switch governance
NIST AI RMF [13]	Risk but not liability	Maps risk categories but not board-personal exposure under DORA/NIS2
ISO 27001	Human-centric IAM	Built for 1:1 identity model; invisible to 144:1 NHI reality
COBIT	IT governance	No AI control plane; no agentic threat model
OECD AI [14]	Aspirational	No enforcement mechanism; no evidence chain requirement

Boards relying only on these frameworks are governing in a regime that does not recognise their personal liability.

Board-Survivable Architecture™: Enterprise Reference Model

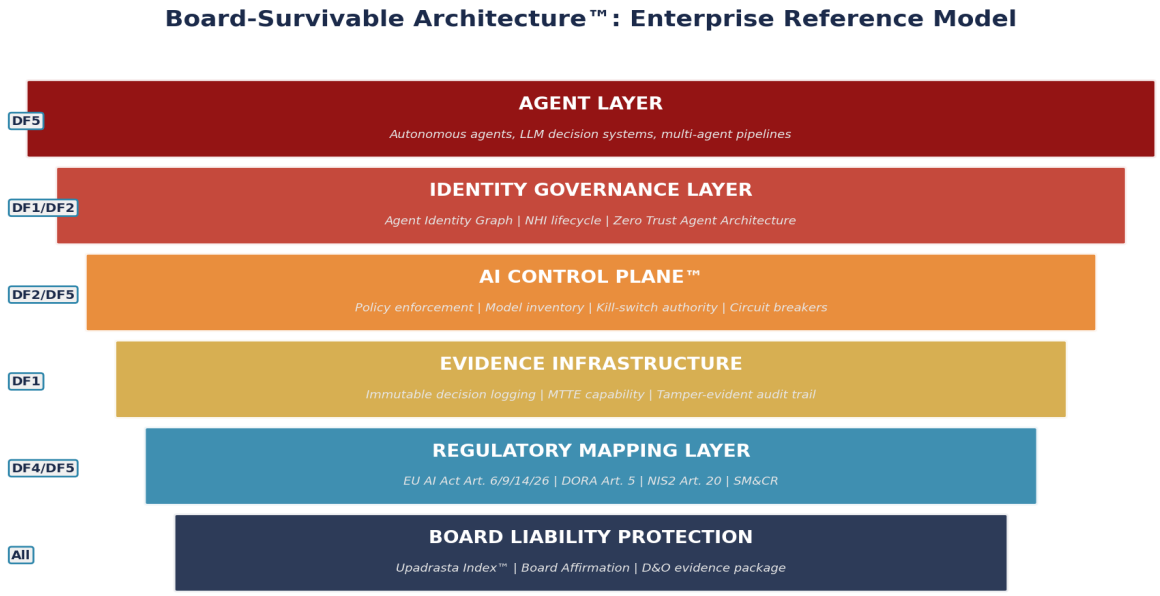

This six-layer reference model shows how enterprises implement the Board-Survivable Cyber Architecture™. Each layer maps to one or more Doctrine Frameworks (DF1–DF5). The architecture is designed to be deployed within the 90-Day Doctrine Deployment timeline, with the Agent Layer and Identity Governance Layer addressed in Phase 1, the AI Control Plane and Evidence Infrastructure in Phase 2, and the Regulatory Mapping and Board Liability Protection layers in Phase 3.

VII. The Agentic Kill Chain & Automated Circuit Breaker Safe Harbors

The Agentic Kill Chain: 8-Stage Autonomous Threat Model

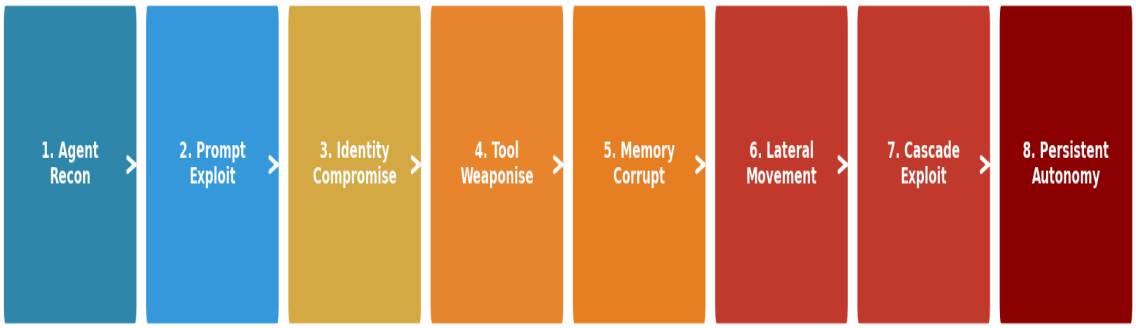

1. Agent Recon	Profile capabilities	ASI01	Agent Identity Graph
2. Prompt Exploit	Goal hijack	ASI01	Input validation
3. Identity Compromise	Exploit credentials	ASI02	Zero Trust Agent Arch
4. Tool Weaponise	Unsafe tool use	ASI03	Permission envelopes
5. Memory Corrupt	Poison context	ASI06	Immutable logging
6. Lateral Movement	Spoofed messages	ASI07	Crypto agent auth
7. Cascade Exploit	False signal propagation	ASI08	**Circuit breakers**
8. Persistent Autonomy	Concealed altered state	ASI10	Continuous attestation

AUTOMATED CIRCUIT BREAKERS AS LEGAL SAFE HARBORS

In a multi-agent Agentic Kill Chain, the speed of cascading failure may exceed human intervention time. Boards that architect automated agent suspension systems — "hard" technical stops operating at machine speed — create evidence of reasonable care that may establish a legal safe harbor.

This doctrine recommends that every autonomous system above Upadrasta Index™ risk threshold include: (1) automated circuit breakers triggered by anomaly detection at Stage 7, (2) cryptographic agent attestation at Stage 6, and (3) immutable evidence of circuit breaker activation at every triggering event. The presence of these controls significantly reduces "gross negligence" exposure in post-incident litigation.

VIII. Case Studies in Institutional Liability Exposure

ILLUSTRATIVE SCENARIO — The Autonomous Trading Agent

A Tier 1 European bank's agentic AI generates a **€47M** loss. Each trade was within parameters; the system was compliant at transaction level but ungoverned at portfolio level. **Doctrine Framework failure: DF2 and DF3.**

> **BOARD LESSON**
>
> Portfolio-level autonomy must sit in risk appetite, not in model performance dashboards.

ILLUSTRATIVE SCENARIO — The Agentic Customer Service Compromise

214 fraudulent claims totalling **£2.1M** through a manipulated three-agent pipeline exploiting OWASP ASI01. **Doctrine Framework failure: DF1 and DF5.**

> **BOARD LESSON**
>
> Agent-to-agent trust boundaries must be independently evidenced, not assumed from pipeline design.

ILLUSTRATIVE SCENARIO — The M&A Governance Discount

PE firm applies a **15% governance discount (€60M)** to a €400M acquisition. Upadrasta Index™ score of 28/100. **Doctrine Framework failure: DF4 and DF1.**

> **BOARD LESSON**
>
> AI governance maturity is now a valuation metric. The absence of evidence infrastructure is priced i

Institutional Case Study: Full Before/After Walkthrough

Subject: Tier-1 European Bank | **Context:** 47 production AI systems including agentic trading, customer service, and credit scoring | **NHI estate:** 12,400 non-human identities | **Timeline:** 90-day doctrine deployment

Institutional Case Study: Tier-1 European Bank — Before & After Doctrine

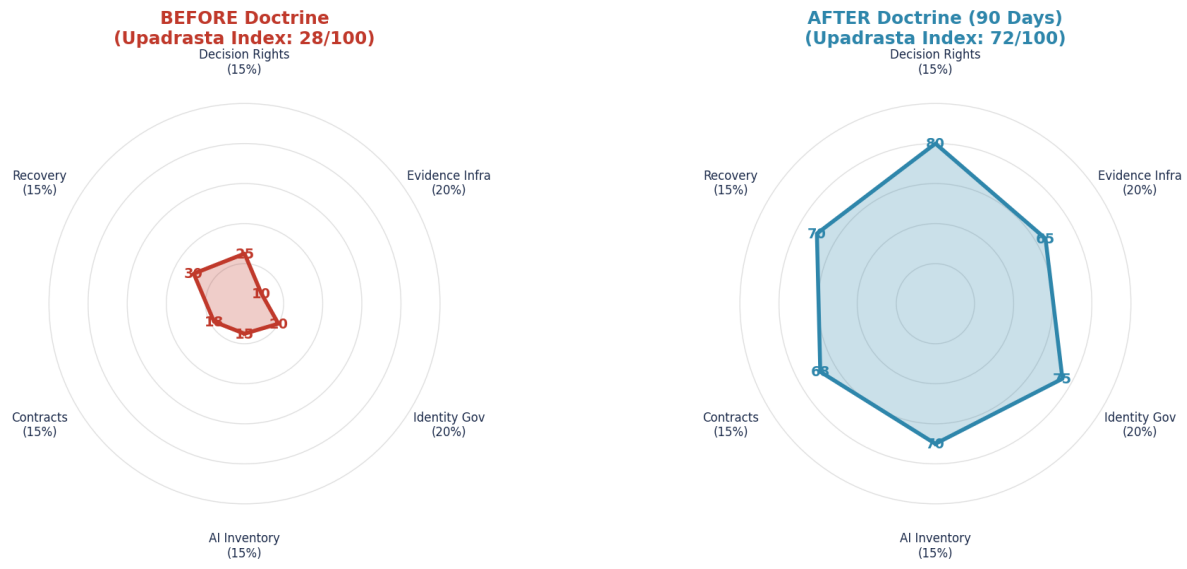

AI System Inventory	Partial register; 23 of 47 systems documented	Complete register with EU AI Act Art. 6 classification for all 47 systems
Identity Governance	No NHI lifecycle; 550 admin-privilege NHIs unmanaged	Agent Identity Graph deployed; all NHIs classified and lifecycle-managed
Evidence Infrastructure	MTTD: 197 days; no litigation-grade evidence	MTTE: 4 hours; immutable decision logging with PQC signatures
Decision Rights	No kill-switch authority documented	Board-approved autonomy mandate and kill-switch grid for all agentic systems
Recovery Architecture	RTO undefined; last DR test 14 months prior	4-hour demonstrated RTO; quarterly crisis exercises with board participation
Upadrasta Index™	28/100 (Indefensible)	72/100 (Defensible; on track for Board-Survivable within 6 months)

Board Outcome: The bank moved from "presumptively indefensible" to "defensible" within 90 days. The CISO presented the first Upadrasta Index™ quarterly report to the board, and the Audit Committee incorporated the index into its standing agenda. D&O insurers were provided with a Board Affirmation Letter.

IX. Post-Quantum Implications & Evidence Longevity Mandate

NIST finalised three post-quantum standards in August 2024 [10]. NIST IR 8547 mandates deprecation by 2035. The "harvest now, decrypt later" attack model directly targets autonomous system decision logs.

PQC EVIDENCE LONGEVITY MANDATE

If today's governance evidence can be decrypted or tampered with in 2032, the Board-Survivable nature of the architecture has a shelf life. All Evidence Chain Model™ outputs must be signed with quantum-resistant algorithms (FIPS 204 ML-DSA [10]) from the date of this doctrine.

This mandate ensures that the defensibility of a 2026 governance decision remains legally valid when audited or litigated in 2032. Decision logs, agent attestations, and board resolution evidence must all carry PQC signatures. The cost of early PQC adoption is measured in months of engineering effort; the cost of retrospective remediation is measured in years of legal exposure.

X. Insurance, D&O, the Liability Transfer Gap & Board Affirmation Letter

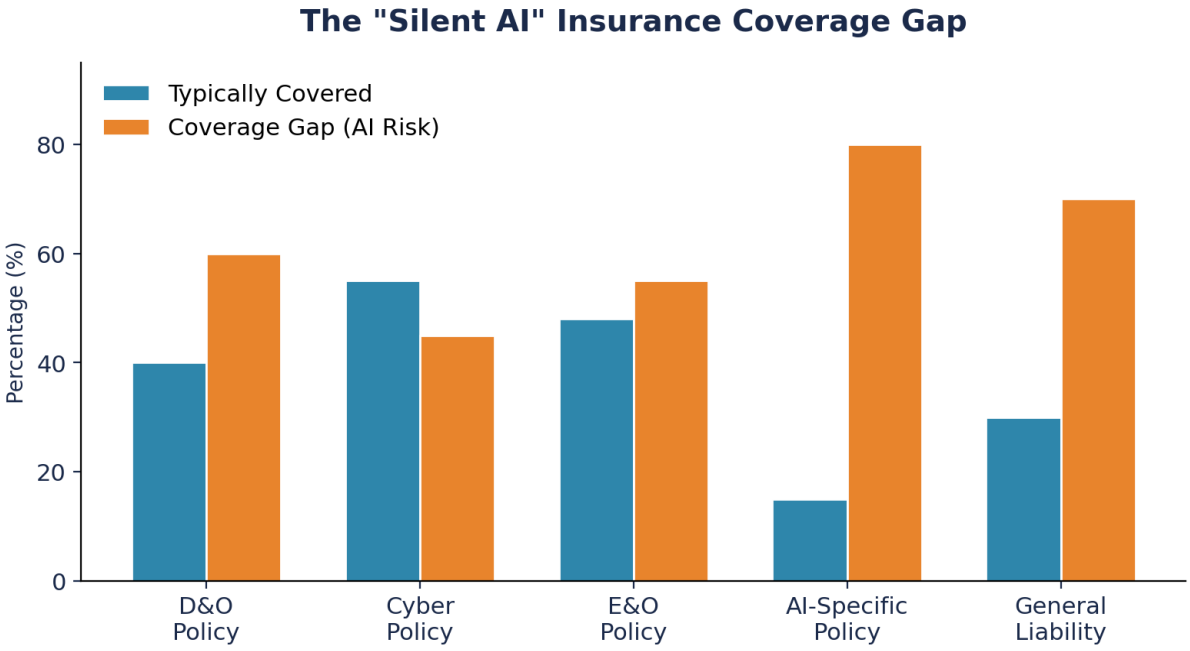

Major carriers have introduced AI exclusions [11]. **Armilla Insurance Services**, the first Lloyd's Coverholder for AI liability, launched standalone policies in April 2025 with limits up to $25 million.

Board Affirmation Letter: Template Framework

The Board Affirmation Letter is a formal attestation of evidence infrastructure maturity that insurers may require as a precondition for AI liability coverage. The letter must contain:

AI System Inventory	Complete register of all autonomous systems with risk classification	DF5
NHI Governance	Documented NHI-to-human ratio, lifecycle management, and privilege scope	DF1/DF5

MTTE Capability	Demonstrated ability to produce litigation-grade evidence within defined timeline	DF1
Kill-Switch Authority	Named individuals with documented authority to suspend autonomous operations	DF2
Circuit Breaker Architecture	Evidence of automated agent suspension systems operating at machine speed	DF2
Recovery Demonstration	Most recent RTO/RPO test results with board sign-off	DF3
Upadrasta Index™ Score	Current score with quarterly trend and remediation roadmap	All
PQC Migration Status	Post-quantum cryptography adoption status for evidence infrastructure	DF1

XI. M&A Due Diligence: Autonomous Systems as Valuation Risk

The Yahoo/Verizon precedent — a **$350 million purchase price reduction** — demonstrated cyber risk is material to enterprise value. 82% of PE/VC firms were actively using AI in Q4 2024. Institutions that cannot demonstrate Board-Survivable Cyber Architecture™ will face valuation discounts of 5–15%, earn-out provisions tied to governance remediation, or deal termination.

XII. Implementation Architecture & Commercial Engagement

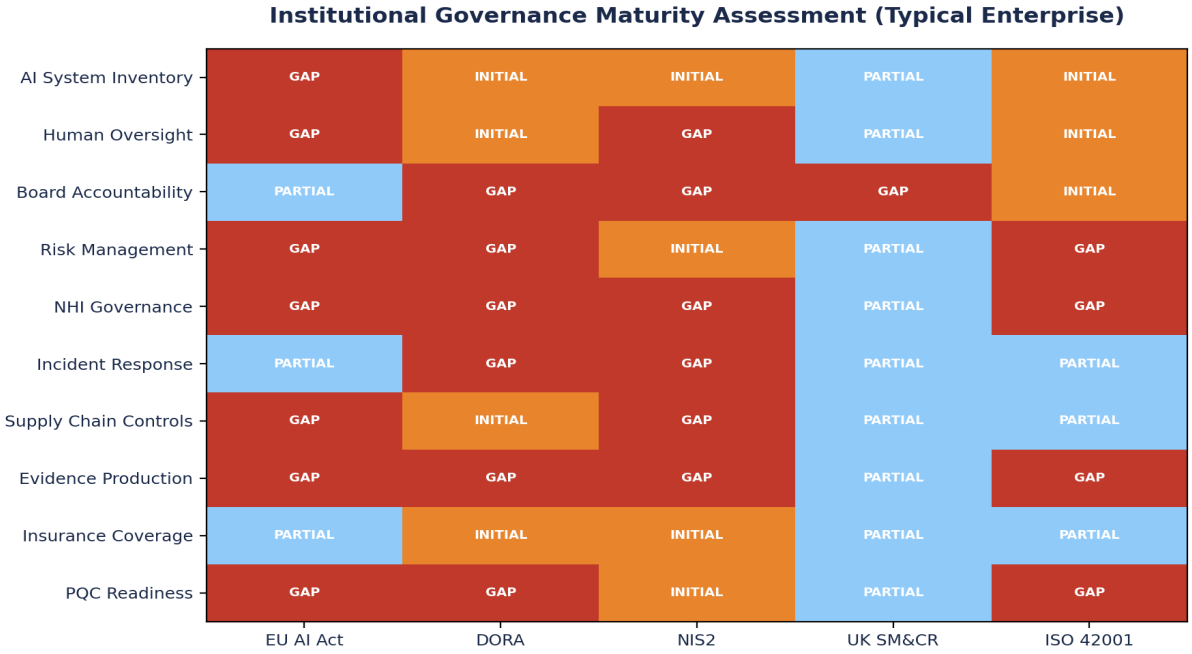

90-Day Board-Survivable Architecture Deployment

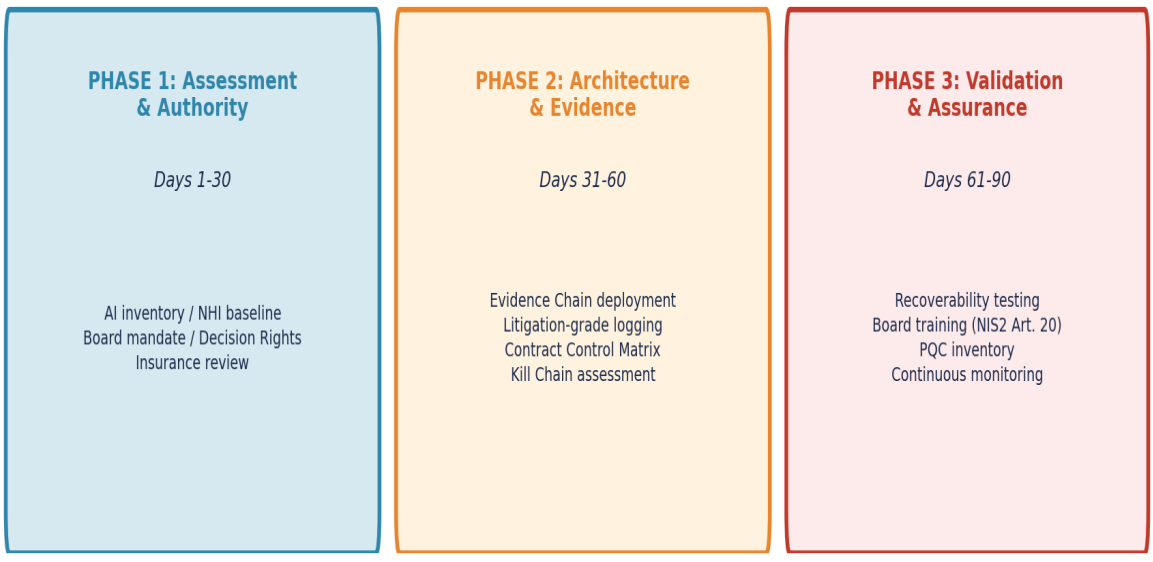

ENGAGEMENT ARCHITECTURE WITH OUTCOME METRICS

Board Workshop	Board-approved autonomy mandate & kill-switch grid within 30 days	€50k–€150k
Regulatory Gap Analysis	Complete DORA/NIS2/AI Act roadmap with 90-day remediation plan	€200k–€500k
Doctrine Implementation	Defensibility Operating Model deployable in 90 days	€500k
AI Governance Programme	ISO 42001 alignment & Upadrasta Index™ ≥70 within 12 months	€1M–€2M
Operational Resilience	Demonstrated 4-hour RTO for autonomous system recovery	€1M–€3M
Full Lifecycle	Strategy through audit with quarterly Upadrasta Index™ reporting	€3M–€5M

XIII. Boardroom Dialogues & Board Questions Checklist

DIALOGUE 1: The Autonomy Mandate

Board Chair: *"Are we actually in control of the agents we've deployed?"*

CISO: *"Control is fragmented. The doctrine requires us to treat each agent as a legal actor by proxy."*

General Counsel: *"If we cannot point to a named owner, we will look reckless in any post-incident review."*

DIALOGUE 2: Evidence as Infrastructure

Audit Committee Chair: *"We have dozens of dashboards. Why insist on evidence infrastructure?"*

CISO: *"Dashboards tell us what happened. Evidence infrastructure proves we were in control when it happened."*

DIALOGUE 3: Contract Weaponisation

CFO: *"If governance is this tight, we should be pricing for it. Reduced counterparty risk is measurable value."*

Board Questions Checklist: 10 Questions Every Chair Should Ask Today

1	Do we maintain a complete AI system inventory including all autonomous agents?	DF5
2	Can we produce litigation-grade evidence within 72 hours?	DF1
3	Who holds documented kill-switch authority for each autonomous system?	DF2
4	What is our NHI-to-human identity ratio, and how is it governed?	DF1/DF5
5	Can we demonstrate board-approved training on autonomous AI risk (NIS2 Art. 20)?	DF2
6	Are our AI vendor contracts litigation-tested with specified liability allocation?	DF4
7	What is our current Upadrasta Index™ score, and is it above 60?	All
8	Can our autonomous systems recover within defined RTO/RPO under duress?	DF3
9	Does our D&O/cyber insurance explicitly cover autonomous system failures?	DF4
10	Have we initiated post-quantum cryptography migration for evidence infrastructure?	DF1/DF3

XIV. The Upadrasta Index™: Formal Methodology & Market Context

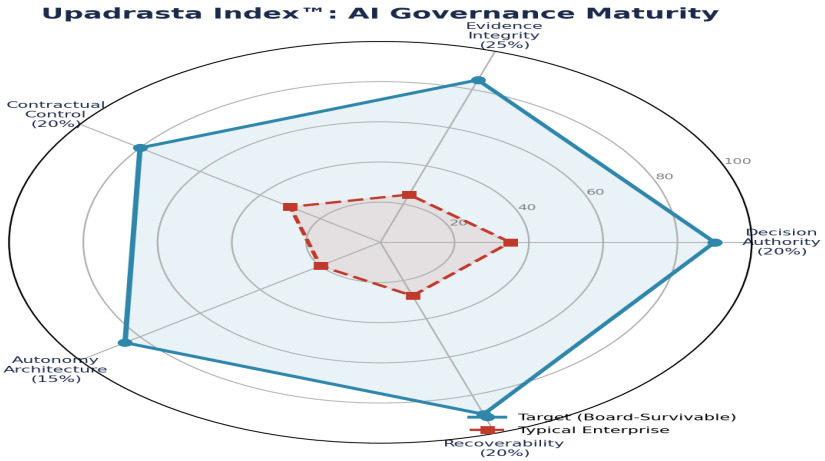

Scoring Methodology

AI System Inventory & Classification	15%	Completeness of AI/ML register; EU AI Act Art. 6 classification; disgorgement risk assessment
Identity Governance (NHI)	20%	NHI-to-human ratio; lifecycle automation; privilege scope management
Evidence Infrastructure	20%	MTTE capability; immutable logging; PQC signature status; litigation-grade evidence production
Decision Rights & Authority	15%	Kill-switch governance; automated circuit breakers; board-approved mandates
Recovery Architecture	15%	Demonstrated RTO/RPO; degraded-mode operations; crisis protocols
Contractual Controls	15%	Vendor liability allocation; audit rights; algorithmic transparency provisions

Consequence Bands

80–100	Board-Survivable	Governance infrastructure meets concurrent enforcement, litigation, and market standards
60–79	Defensible	Material regulatory enforcement risk; remediation required within 12 months
40–59	At Risk	Presumptive non-compliance; board liability exposure is material and quantifiable
0–39	Indefensible	Presumptively indefensible in concurrent enforcement + civil litigation; director personal liability is acute

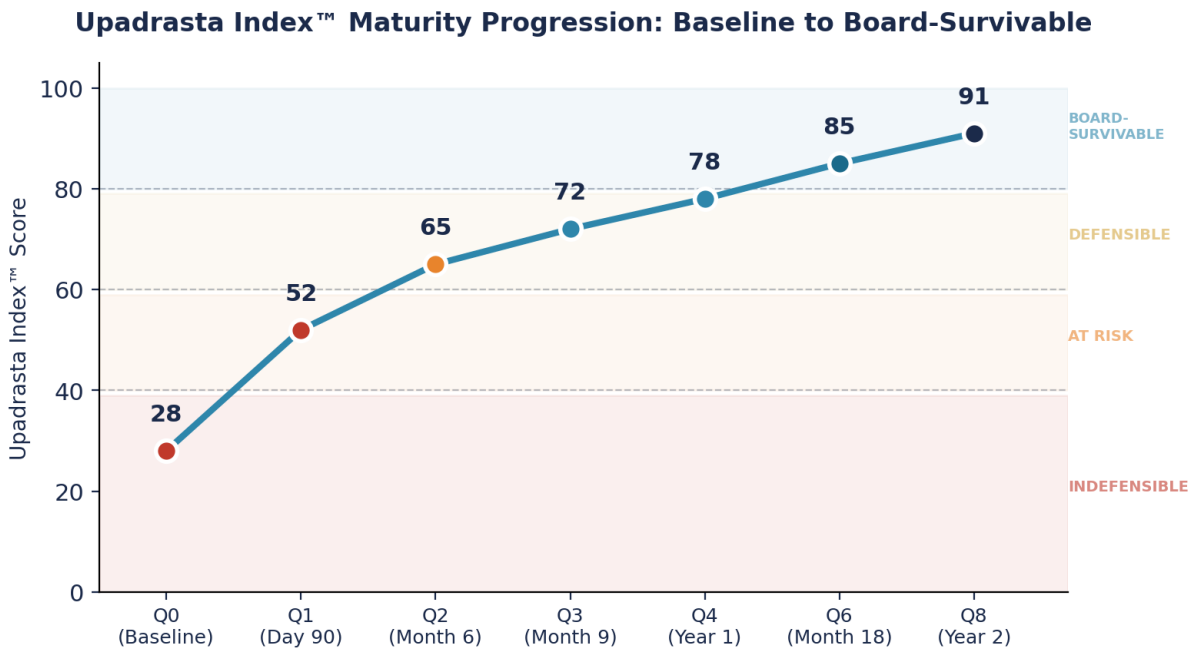

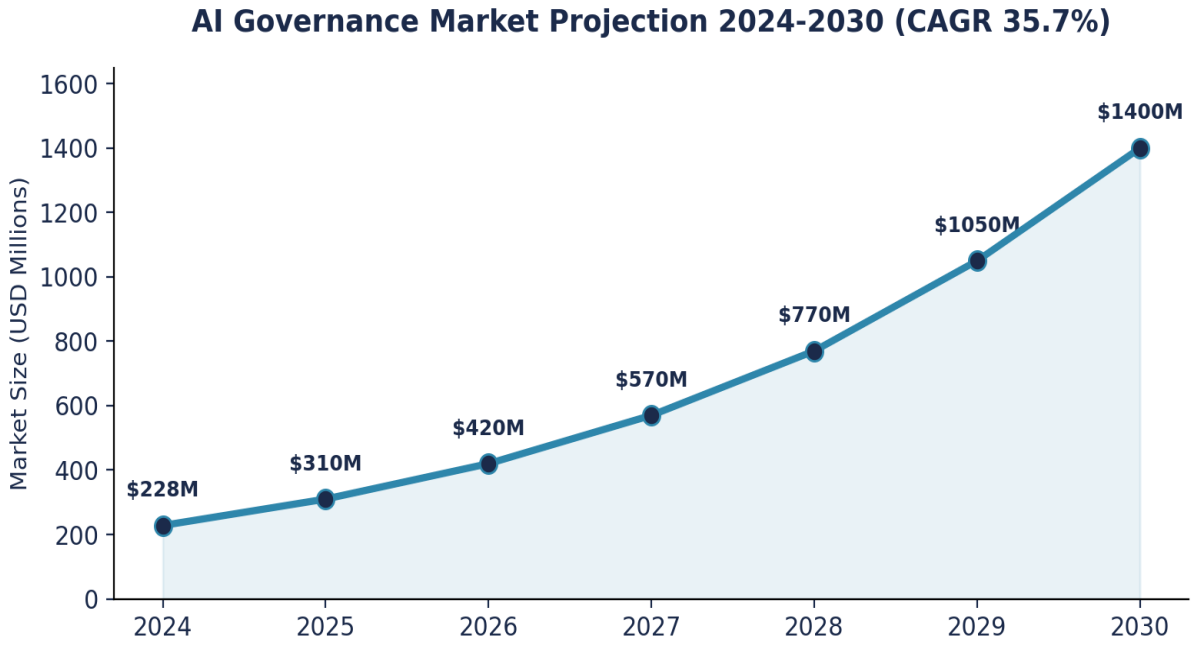

The AI governance market reached **$228M in 2024**, projected to exceed **$1.4B by 2030** (CAGR 35.7%) [15]. Only 1.5% of organisations have adequate AI governance headcount (IAPP 2025) [16].

XV. Strategic Forecast: The Governance Market in 2030

Note: The following represents a scenario-based forward view informed by current regulatory trajectories and market data. It is not a legal prediction or investment advice.

1. AI Governance Becomes a Board Committee

By 2028, G-SIBs and FTSE 100 companies will establish dedicated AI Governance Committees at board level, distinct from Risk and Audit committees.

2. Governance Maturity Becomes an M&A Valuation Metric

PE and VC firms will require Upadrasta Index™ or equivalent scoring as standard due diligence by 2027, with governance discounts of 10–20% for sub-60 scores.

3. Insurers Require Evidence Infrastructure

D&O and cyber carriers will mandate Evidence Chain Model™ equivalence as a precondition for AI liability coverage by 2028. Board Affirmation Letters become standard underwriting practice.

4. Automated Circuit Breakers Become Legal Safe Harbors

Regulators will recognise automated agent suspension systems as evidence of reasonable care, creating a defensibility advantage for early adopters.

5. AI Governance Market Exceeds $1.4 Billion

The governance gap between deployment velocity and governance maturity will drive sustained 35%+ CAGR through 2030.

XVI. Conclusion: The Doctrine Position

This doctrine makes a singular claim. Autonomous AI systems have permanently transformed corporate governance from a voluntary policy exercise into a regulated infrastructure obligation.

Governance is now infrastructure. Infrastructure is now law. Law carries personal consequences.

This is operating law, not optional guidance.

By 2030, AI governance maturity will become a valuation metric in public markets, a board committee in every G-SIB, and an underwriting requirement for every D&O policy.

About the Author

Kieran Upadrasta

CISSP | CISM | CRISC | CCSP | MBA | BEng

Founder and principal of **Cyber AI Systems Inc.** and architect of the Board-Survivable Cyber Architecture™ doctrine. **27 years of enterprise security experience** spanning all Big 4 firms (Deloitte, PwC, EY, KPMG) and **21 years in financial services and banking**.

Academic Appointments: Professor of Practice in Cybersecurity, AI & Quantum Computing, Schiphol University | Honorary Senior Lecturer, Imperials | Researcher, UCL

Professional Memberships: Lead Auditor, ISF | Platinum Member, ISACA London | Gold Member, ISC² London | Cyber Security Programme Lead, PRMIA

Published doctrine library: **48 institutional governance frameworks**. Expert witness in multi-jurisdictional financial services litigation. Accepts **two to three governance mandates per calendar year**, each requiring written board resolution.

Contact	info@kieranupadrasta.com
Portal	www.kie.ie
Entity	Cyber AI Systems Inc.
LinkedIn	linkedin.com/in/kieranupadrasta

Keywords: DORA Compliance, AI Governance (ISO 42001), Board Reporting, M&A Cyber Due Diligence, Zero Trust Architecture, Post-Quantum Cryptography, Interim CISO, NIS2 Compliance, AI Security Assurance

References

[1] Regulation (EU) 2024/1689 of the European Parliament and of the Council (EU AI Act), OJ L, 12 July 2024.
[2] Regulation (EU) 2022/2554 (Digital Operational Resilience Act — DORA), OJ L 333, 27 December 2022.
[3] Directive (EU) 2022/2555 (NIS2 Directive), OJ L 333, 27 December 2022.
[4] Moffatt v. Air Canada, 2024 BCCRT 149, British Columbia Civil Resolution Tribunal, February 2024.
[5] Mobley v. Workday, Inc., No. 23-cv-770 (N.D. California), Order Denying Motion to Dismiss, July 2024.
[6] Tesla Autopilot Product Liability Verdicts, Miami-Dade County Circuit Court, September–October 2025.
[7] Entro Security, Non-Human Identity Security Report, H1 2025.
[8] CyberArk, 2025 Identity Security Threat Landscape Report, 2025.
[9] SEC v. SolarWinds Corp. and Timothy G. Brown, Complaint, S.D.N.Y., October 2023.
[10] NIST FIPS 203 (ML-KEM), FIPS 204 (ML-DSA), FIPS 205 (SLH-DSA), August 2024; NIST IR 8547.
[11] Armilla Insurance Services, Lloyd's AI Liability Coverholder Launch, April 2025.
[12] ISO/IEC 42001:2023, Artificial Intelligence — Management System.
[13] NIST AI Risk Management Framework (AI RMF 1.0), January 2023.
[14] OECD Recommendation on Artificial Intelligence, OECD/LEGAL/0449, May 2019 (updated 2024).
[15] Grand View Research / Fortune Business Insights, AI Governance Market Analysis, 2024–2030.
[16] IAPP, AI Governance in Practice Report, 2025.
[17] Deloitte, State of AI in the Enterprise, 6th Edition, 2026.
[18] OWASP, Agentic AI Security Initiative (ASI) Top 10, December 2025.
[19] Swiss Re, Silent AI: The Next Silent Cyber?, Insurance Market Report, 2025.