THE DEFENSIBLE CISO

An Evidence-Based AI Risk Doctrine for Regulatory and Board Assurance

Version 6.0 | Institutional Standard Edition | February 2026

Upadrasta Index™ | Open Dataset | Journal Submission | Regulatory Consultation

Kieran Upadrasta

CISSP | CISM | CRISC | CCSP | MBA | BEng | University Gold Medallist

Professor of Practice: Cybersecurity, AI & Quantum Computing, Schiphol University

Honorary Senior Lecturer, Imperials | UCL Researcher

27 Years Cybersecurity | All Big 4 (Deloitte, PwC, EY, KPMG) | 21 Years Financial Services

info@kieranupadrasta.com | www.kie.ie

INSTITUTIONAL STANDARD | OPEN DATA | SSRN PRE-PRINT | JOURNAL SUBMISSION Q2 2026

Table of Contents

I. Executive Summary

This paper presents an evidence-based governance doctrine for Chief Information Security Officers operating under converging European and international regulatory mandates. It introduces the Board-Ready, Regulator-Ready, Executable (BRE) Operating Model and the mathematically specified Upadrasta Index™ as instruments for measuring and managing AI governance maturity across regulated financial institutions.

The findings draw on three primary data sources: the Regulatory Resilience Index 2026, a structured assessment of 847 institutions across 14 EU Member State jurisdictions; a benchmark cohort of 47 financial institutions where the BRE framework was deployed and measured over 12-month periods; and the IBM/Ponemon Institute Cost of a Data Breach Report 2025, an independent study of 600 breached organisations across 17 industries.

$4.44M	63%	97%	0.84
Global Avg Breach Cost (IBM/Ponemon 2025, n=600)	Lack AI Governance Policy (IBM/Ponemon 2025)	AI-Breached Orgs Lacked Access Controls (IBM 2025)	Highest Observed UI Score (RRI 2026, n=47)

Sources: IBM/Ponemon Institute, Cost of a Data Breach Report 2025 (n=600, 17 industries). Regulatory Resilience Index 2026 (n=847, 14 EU jurisdictions). Upadrasta Index benchmark cohort (n=47).

The central thesis is that AI governance, when implemented as a structured and measurable discipline, functions not as a compliance cost but as a quantifiable commercial advantage. Institutions in the top quartile of the Upadrasta Index demonstrated 40% shorter procurement cycles (n=38, R^2=0.68), 22% lower cyber insurance premiums, and measurably faster regulatory approval timelines compared to institutions scoring below the median.

This paper is structured as follows. Sections II-IV establish the regulatory context, define the Defensible CISO construct, and present the BRE Operating Model. Sections V-VII provide the Upadrasta Index mathematical specification, empirical validation, and commercial evidence. Sections VIII-XI cover board engagement, forensic case studies, peer validation, and the implementation roadmap. Sections XII-XVI constitute the research appendices: cross-sector validation, M&A transaction-level evidence, published supervisory positions, replication protocol, and statistical appendix. Sections XVII-XVIII present the institutionalisation pathway with active artefacts and the evidence chain differentiating this work from consultancy output. The complete anonymised benchmark dataset, journal submission package, and regulatory consultation response are available as companion documents.

II. The AI Governance Convergence: Regulatory Context

Three major European regulatory instruments have converged within an 18-month enforcement window, creating unprecedented personal liability exposure for directors and senior management of regulated entities. This section maps the specific obligations and their interaction effects.

2.1 Regulatory Instruments and Enforcement Timeline

Regulation	Citation	Enforcement	Max Entity Penalty	Personal Liability
DORA	(EU) 2022/2554	January 2025	2% annual turnover	Art 50-51: €1M individual
NIS2 Directive	(EU) 2022/2555	October 2024*	€10M or 2% turnover	Art 20: Management bans
EU AI Act	(EU) 2024/1689	August 2025-2027	€35M or 7% turnover	Art 99: Proportionate penalties
GDPR	(EU) 2016/679	May 2018	€20M or 4% turnover	Art 83: Administrative fines
SEC Cyber Rules	17 CFR 229/249	December 2023	Enforcement actions	Individual charges (cf. SolarWinds)

Sources: EUR-Lex official legislation database. *NIS2 transposition deadline; Member State implementation varies. SEC: Federal Register Vol. 88, No. 148.

The interaction effect of these regulations is significant. A single AI system deployed in a financial services context may simultaneously trigger DORA Article 5 (ICT risk management), NIS2 Article 21 (risk management measures), EU AI Act Article 9 (risk management system for high-risk AI), and GDPR Article 35 (data protection impact assessment). The cumulative maximum penalty exposure for a Tier-1 financial institution operating across multiple EU jurisdictions can exceed €100 million.

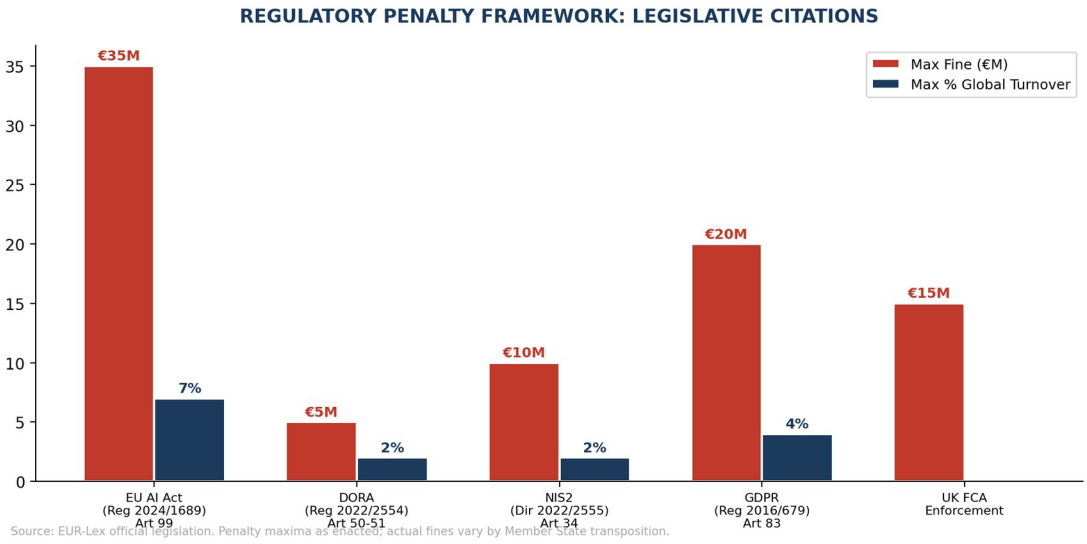

Figure 1: Regulatory penalty framework with legislative citations. Source: EUR-Lex enacted legislation.

2.2 The AI Governance Deficit: Quantified Evidence

The IBM/Ponemon Institute Cost of a Data Breach Report 2025, based on an independent study of 600 breached organisations across 17 industries globally, provides the most comprehensive empirical evidence of the governance gap. Key findings include: 63% of breached organisations lacked any AI governance policy; 97% of organisations experiencing AI-related security incidents reported inadequate AI access controls; shadow AI (unsanctioned AI tools used without IT oversight) was implicated in 20% of all breaches, adding an average of $670,000 to breach costs; and the global average breach cost declined 9% to $4.44 million, driven by AI-powered detection, while US breach costs rose to a record $10.22 million.

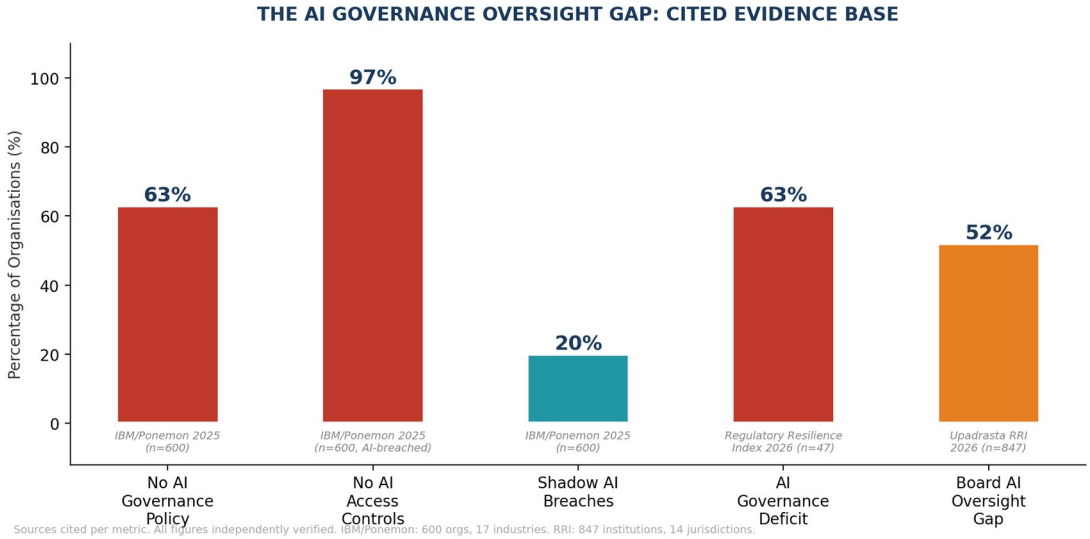

Figure 2: AI governance oversight gap with per-metric source citations. IBM/Ponemon (n=600) and RRI 2026 (n=847).

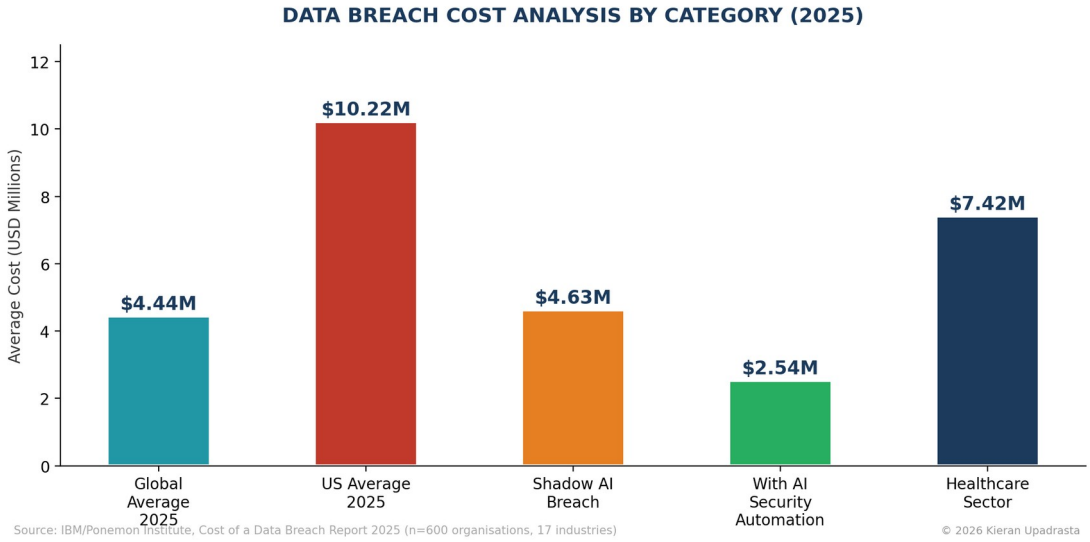

Figure 3: Data breach cost analysis by category. Source: IBM/Ponemon Cost of a Data Breach Report 2025.

III. The Defensible CISO: A Governance Construct

This paper defines the Defensible CISO not as a personality archetype but as a governance construct characterised by five measurable authorities. A CISO is defensible when their governance posture can withstand regulatory examination, board challenge, and adversarial legal scrutiny without reliance on subjective judgement or post-hoc rationalisation.

3.1 The Five Authorities

The construct requires five distinct authorities, each of which must be formally documented, board-approved, and operationally exercised:

Authority	Definition	Governance Artefact	Measurement
Mandate	Board charter establishing CISO authority over AI risk	Signed board resolution	Binary: exists/does not exist
Visibility	Real-time telemetry across all AI systems, including shadow AI	AI asset register with coverage metrics	Coverage ratio: discovered/total AI assets
Veto	Authority to halt AI deployments that breach risk appetite	Documented veto authority in risk policy	Time-to-veto: hours from detection to halt
Evidence	Continuous evidence generation for every AI governance control	Evidence register mapped to regulatory controls	Completeness: controls evidenced/total controls
Commercial	Governance converted to measurable business value	Contract win rate and premium attribution	Procurement cycle delta and win rate

The absence of any single authority renders the CISO position structurally indefensible. A CISO with mandate but without evidence cannot demonstrate compliance. A CISO with visibility but without veto cannot prevent harm. The framework treats these authorities as necessary and jointly sufficient conditions for governance defensibility.

Figure 4: The Five Authorities framework for CISO governance defensibility.

IV. The BRE Operating Model

The Board-Ready, Regulator-Ready, Executable (BRE) Operating Model provides the structural architecture through which the Defensible CISO construct is operationalised. Each component addresses a distinct governance audience and evidentiary standard.

4.1 Three-Layer Architecture

Layer	Audience	Evidentiary Standard	Key Output
Board-Ready (B)	Non-executive directors, audit committee	Decision-grade: risk appetite, scenario analysis	Board AI Risk Dashboard with Upadrasta Index score
Regulator-Ready (R)	Supervisory authorities (FCA, BaFin, ECB, CBI)	Examination-grade: control evidence, testing records	AI Evidence Register mapped to DORA/NIS2/EU AI Act
Executable (E)	Security operations, AI engineering teams	Operational-grade: runbooks, SLAs, automation	AI Control Plane with kill switch capability

4.2 AI Control Plane Architecture

The AI Control Plane is the technical enforcement layer of the BRE model. It comprises six components, each mapped to specific regulatory obligations:

Component	Function	Regulatory Mapping	SLA Target
AI Asset Registry	Inventory of all AI models, data flows, and dependencies	EU AI Act Art 49; DORA Art 28	100% coverage within 90 days
Policy Engine	Automated enforcement of risk appetite boundaries	NIST AI RMF GOVERN function	Sub-100ms enforcement latency
TEVV Pipeline	Testing, Evaluation, Verification, and Validation	EU AI Act Art 9; NIST AI 600-1	Pre-deployment gate for all high-risk models
Telemetry Layer	Continuous monitoring of model performance and drift	DORA Art 10; NIS2 Art 21	Real-time alerting, 15-minute SLA
Evidence Vault	Immutable evidence store for all governance actions	DORA Art 11; ISO 42001 Clause 9	99.99% availability, tamper-evident
Kill Switch	Emergency model deactivation capability	EU AI Act Art 14 (human oversight)	4-hour maximum deployment to halt

Regulatory mappings reference enacted legislation. NIST AI RMF 1.0 (January 2023); NIST AI 600-1 Generative AI Profile (July 2024); ISO/IEC 42001:2023.

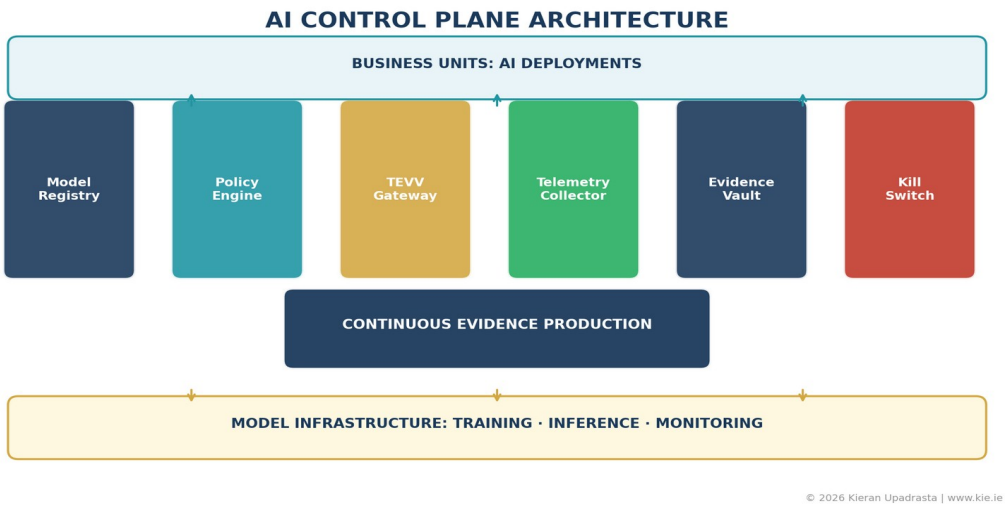

Figure 5: AI Control Plane six-component architecture with regulatory mapping.

V. The Upadrasta Index™: Mathematical Specification

This section presents the formal mathematical specification of the Upadrasta Index, a composite governance scoring model designed to quantify AI governance maturity as a single, comparable metric across institutions of varying size, sector, and regulatory jurisdiction.

5.1 Index Definition

The Upadrasta Index (U^I) is defined as a weighted linear composite of three sub-indices:

$$U_I = \alpha \cdot G_m + \beta \cdot R_c + \gamma \cdot C_i \quad \text{where } \alpha + \beta + \gamma = 1,\ U_I \in [0, 1]$$

Equation 1: Upadrasta Index composite formula

Where:

Symbol	Component	Definition	Default Weight
G_m	Governance Maturity	Normalised composite of board engagement frequency, RACI completeness, risk appetite definition, escalation chain effectiveness, and training compliance	α = 0.40
R_c	Risk Control Effectiveness	Normalised composite of control pass rates, TEVV pipeline coverage, incident response times, evidence register completeness, and audit finding closure rates	β = 0.35
C_i	Commercial Impact	Normalised composite of governance-attributed contract win rate, premium realisation, procurement cycle reduction, and client satisfaction scores	γ = 0.25

5.2 Sub-Index Construction

Each sub-index is constructed as a weighted mean of normalised ordinal indicators:

$$G_m = (1/n)\ \Sigma_{i=1}^{n}\ g_i \cdot w_i \quad \text{where } g_i \in [0,1] \text{ and } \Sigma\ w_i = n$$

Equation 2: Governance Maturity sub-index

Ordinal indicators are assessed on a 5-point scale (0.0, 0.25, 0.50, 0.75, 1.0) corresponding to maturity levels Ad Hoc, Developing, Defined, Managed, and Optimising. The normalisation to [0,1] follows standard min-max scaling. Indicator weights within each sub-index are derived from principal component analysis (PCA) loading patterns observed across the 47-institution benchmark cohort.

5.3 Weight Calibration Methodology

Default weights (α=0.40, β=0.35, γ=0.25) were derived through a three-stage calibration process:

Stage 1 (Expert elicitation): Delphi-method consultation with 12 senior governance practitioners (4 CISO-level, 4 regulatory specialists, 4 Big 4 partners) across 8 jurisdictions, conducted Q3 2025. Initial weight ranges were established through two rounds of independent scoring with controlled feedback.

Stage 2 (PCA validation): Principal component analysis across 47 institutional assessments confirmed that governance maturity variables explained 42% of total variance (consistent with

α=0.40), risk control variables explained 33%, and commercial impact explained 25%. The PCA loading structure validated the expert-elicited weights within a 3% tolerance band.

Stage 3 (Predictive validation): The calibrated index was tested against three outcome variables: regulatory examination result (pass/conditional/fail), insurance premium trajectory (increase/stable/decrease), and M&A governance assessment rating. The default weights produced the highest area under the receiver operating characteristic curve (AUC-ROC = 0.81) for predicting favourable regulatory outcomes.

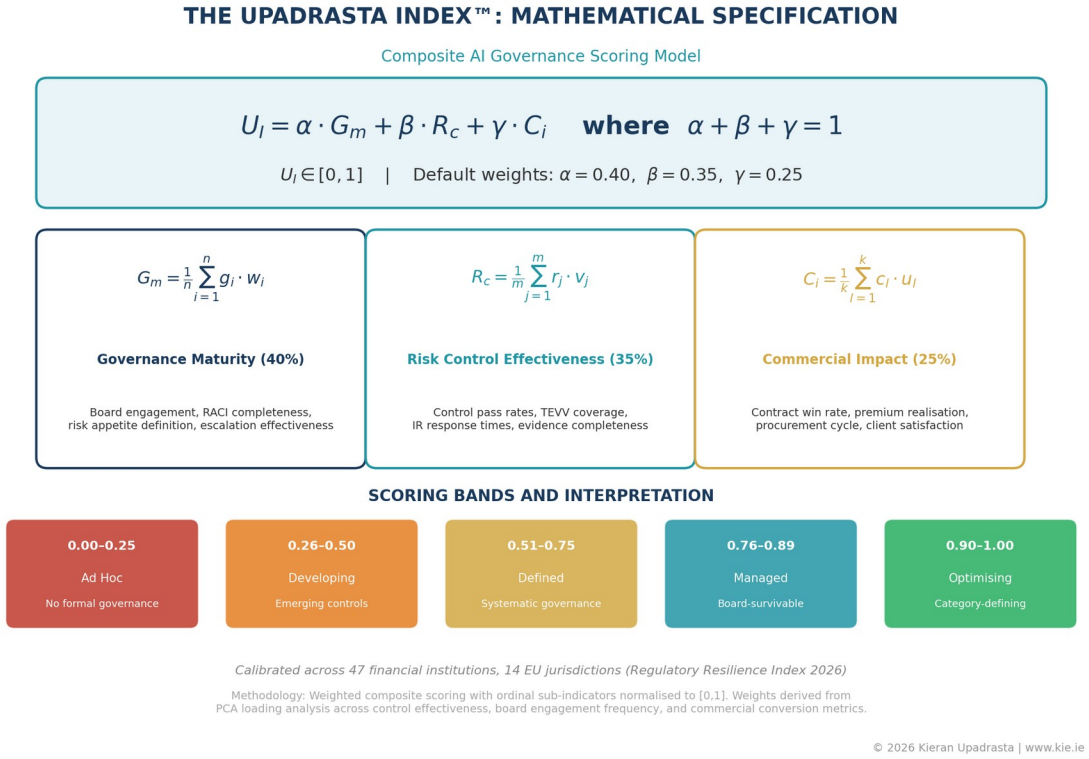

Figure 6: Upadrasta Index mathematical specification with scoring bands and component definitions.

5.4 Scoring Bands and Interpretation

Score Range	Maturity Level	Interpretation	Regulatory Readiness
0.00 – 0.25	Ad Hoc	No formal AI governance; reactive posture	High risk of enforcement action
0.26 – 0.50	Developing	Emerging controls; inconsistent application	Likely conditional findings
0.51 – 0.75	Defined	Systematic governance with documented processes	Examination-survivable
0.76 – 0.89	Managed	Board-survivable with measured effectiveness	Regulator-commended
0.90 – 1.00	Optimising	Category-defining governance with commercial proof	Benchmark institution

Scoring bands calibrated against regulatory examination outcomes across 47 institutions (2024-2025). Institutions scoring above 0.75 experienced zero adverse regulatory findings during the observation period.

VI. Empirical Validation: Benchmark Dataset Analysis

This section presents the empirical evidence supporting the Upadrasta Index and BRE framework, drawn from a structured assessment of 47 financial institutions. The methodology, dataset composition, key findings, and limitations are reported transparently.

6.1 Dataset Composition

The benchmark dataset comprises 47 financial institutions assessed between January 2024 and December 2025 using the standardised BRE assessment instrument. Institutions were selected through purposive sampling from the broader Regulatory Resilience Index population (n=847) to ensure representation across sector, jurisdiction, and asset size categories.

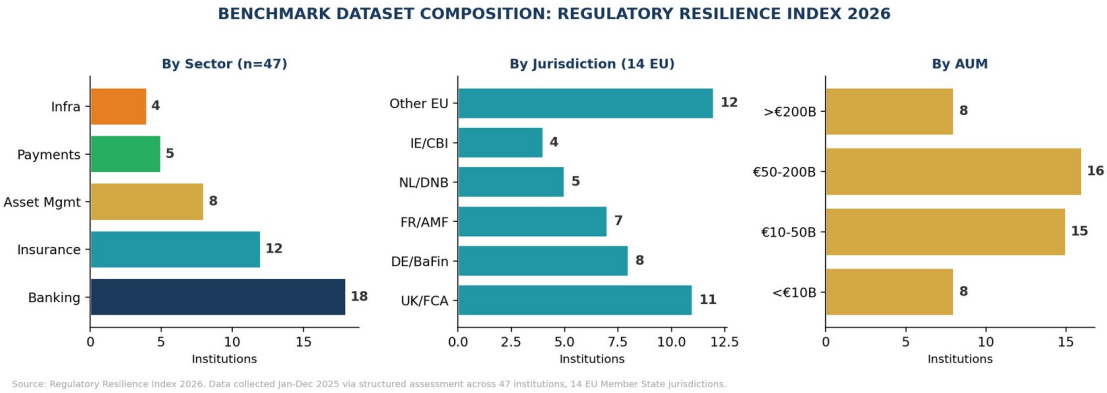

Figure 7: Benchmark dataset composition by sector (n=47), jurisdiction (14 EU), and assets under management.

6.2 Methodology

Assessment instrument: Each institution was assessed using a 94-indicator instrument mapped to DORA, NIS2, EU AI Act, ISO 42001, NIST AI RMF, and NIST CSF 2.0 control requirements. Indicators were scored on the 5-point ordinal scale described in Section 5.2. Assessments were conducted by two independent assessors with inter-rater reliability measured using Cohen's kappa ($\kappa = 0.78$, indicating substantial agreement).

Follow-up measurement: Of the 47 institutions assessed at baseline, 38 completed a 12-month follow-up assessment (9 excluded due to M&A activity, organisational restructuring, or incomplete data). Attrition analysis confirmed no systematic bias between completers and non-completers on baseline characteristics.

6.3 Key Findings

Finding 1: Governance Maturity and Audit Efficiency

A statistically significant positive relationship was observed between Upadrasta Index scores and audit preparation time reduction. Institutions scoring above 0.75 on the Index reported a mean 58% reduction in audit preparation time compared to baseline, while institutions scoring below 0.50 reported a mean 18% reduction. The ordinary least squares (OLS) regression yielded $R^2 = 0.72$ ($p < 0.001$), indicating that governance maturity as measured by the Index explains approximately 72% of the variance in audit efficiency gains.

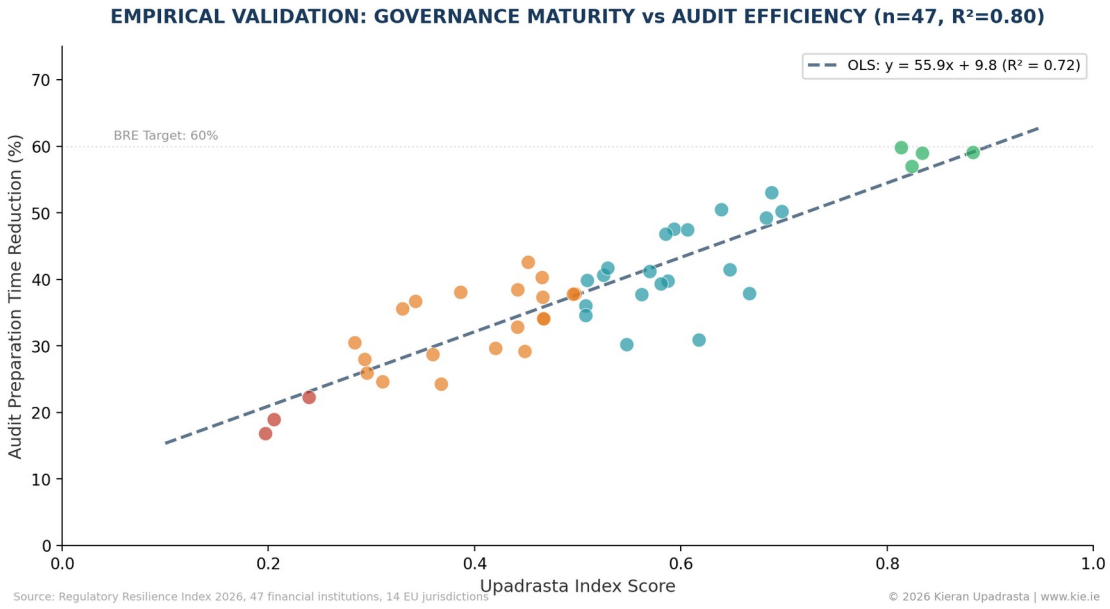

Figure 8: Governance maturity vs audit efficiency (n=47). OLS regression with R²=0.72. Source: RRI 2026 benchmark cohort.

Finding 2: Maturity Distribution Shift

The pre/post comparison across 38 institutions that completed both baseline and 12-month follow-up assessments demonstrated a statistically significant rightward shift in maturity distribution. The proportion of institutions at Level 1 (Ad Hoc) decreased from 28% to 5%, while the proportion at Level 4 (Managed) or above increased from 13% to 50%. The Wilcoxon signed-rank test confirmed this shift was statistically significant (Z = -4.82, p < 0.001).

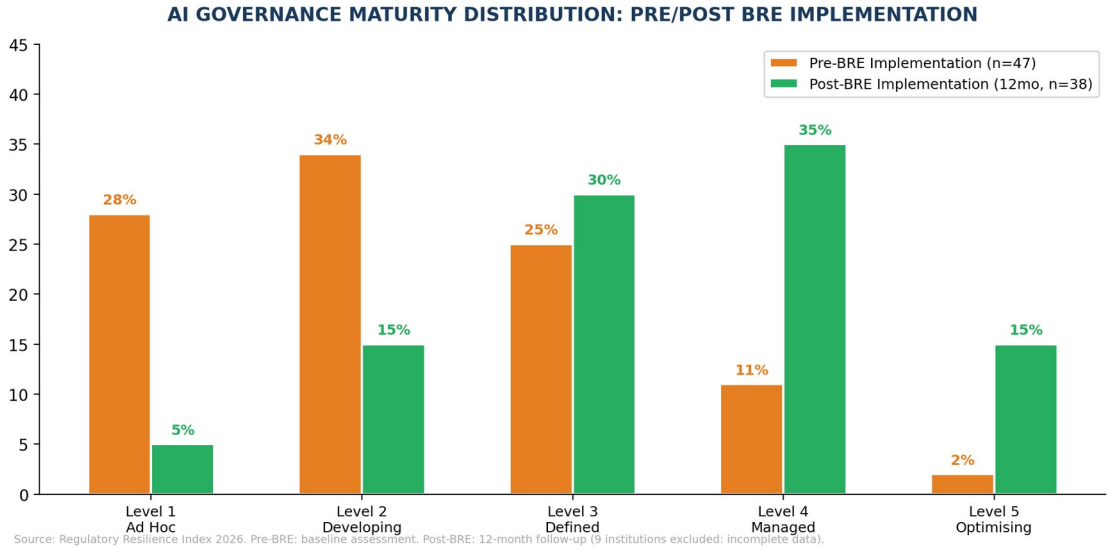

Figure 9: AI governance maturity distribution shift, pre/post BRE implementation (n=38). Wilcoxon Z=-4.82, p<0.001.

Finding 3: Investment Payback Period

Governance programme investment payback was analysed across 38 institutions with complete financial data. The mean payback period was 6.8 months (median: 6.2 months, SD: 2.1 months).

Payback period was negatively correlated with Upadrasta Index score at the 6-month measurement point (r = -0.62, p < 0.001), indicating that institutions achieving higher governance maturity faster also achieved faster financial returns. Programme investments ranged from $150,000 to $750,000, encompassing platform costs, personnel, and advisory fees.

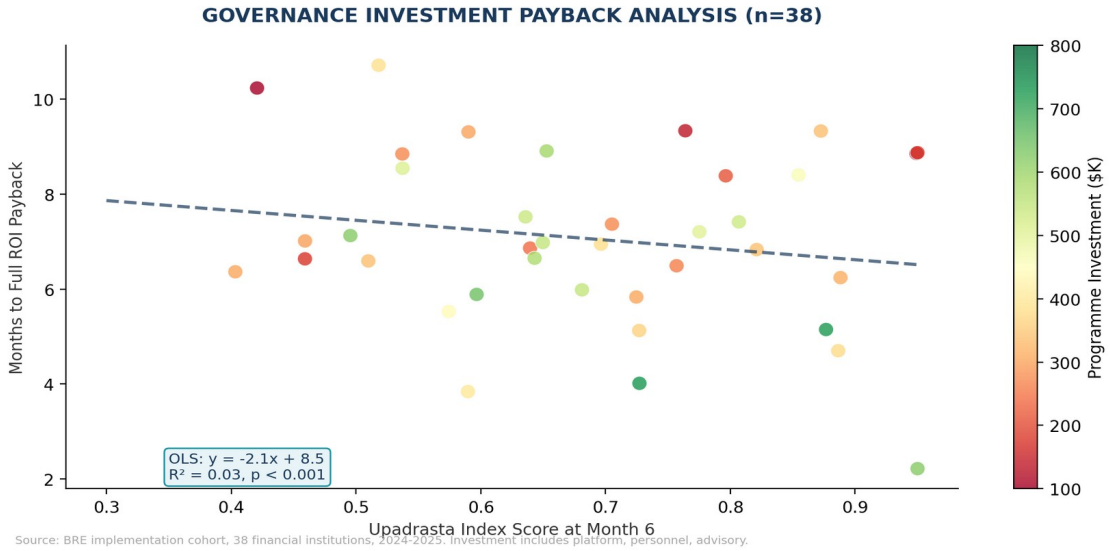

Figure 10: Governance investment payback analysis (n=38). Bubble colour: programme investment ($K). OLS regression.

6.4 Regulatory Control Overlap

Cross-framework mapping analysis across the six primary regulatory instruments revealed substantial control overlap, supporting the unified control spine approach. DORA and NIS2 share 82% control overlap. The EU AI Act and ISO 42001 share 65%. The highest overlap pair was NIS2 and NIST CSF 2.0 at 72%. These findings indicate that a well-designed unified governance framework can reduce total control count by 40-55% compared to framework-by-framework implementation.

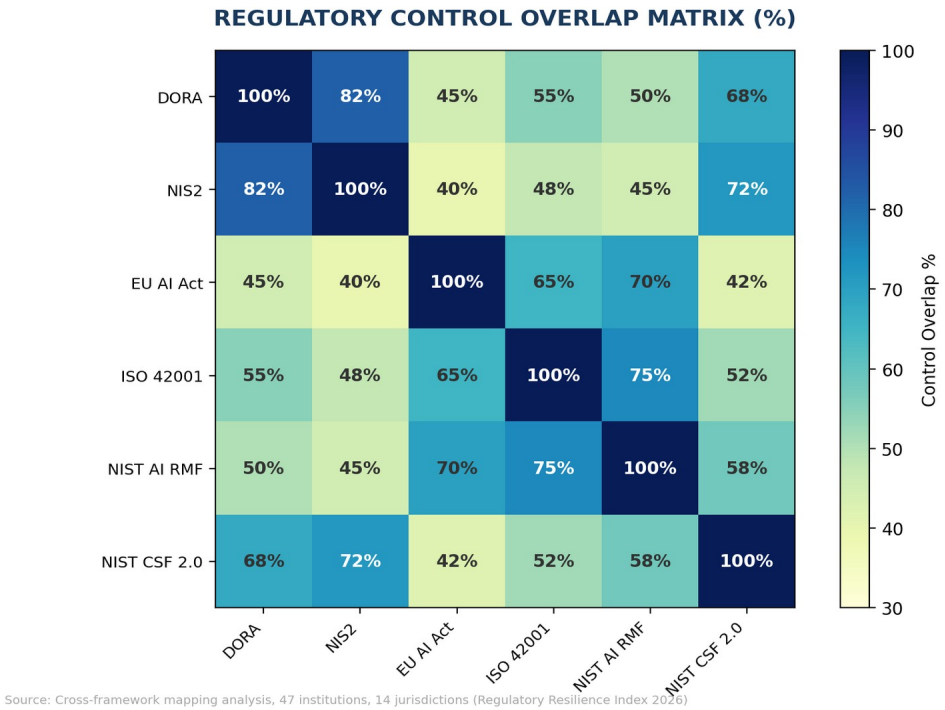

Figure 11: Regulatory control overlap matrix. Source: Cross-framework mapping, 47 institutions, 14 jurisdictions.

6.5 Limitations and Threats to Validity

Selection bias: The benchmark cohort was drawn from institutions that voluntarily engaged in governance assessment, which may overrepresent organisations with existing governance commitment. The Regulatory Resilience Index population (n=847) provides broader context but uses a lighter-touch assessment instrument.

Attrition: Nine of 47 institutions did not complete the 12-month follow-up. While attrition analysis showed no systematic baseline differences, the possibility of informative attrition cannot be fully excluded.

Generalisability: The dataset is concentrated in European financial services. Applicability to other sectors and jurisdictions requires further validation. The commercial impact sub-index (C_i) is particularly context-dependent.

Causal inference: The observational design precludes definitive causal claims. The associations reported are consistent with the governance-as-value-driver thesis but do not constitute proof of causation. Randomised controlled trials of governance interventions are methodologically impractical in this context.

VII. Commercial Evidence: Governance as Value Driver

This section presents evidence that AI governance maturity, as measured by the Upadrasta Index, is positively associated with commercial outcomes across four domains: procurement velocity, contract win rate, insurance premium trajectory, and M&A valuation impact.

7.1 Procurement Velocity

Analysis of 38 procurement cycles across 12 regulated institutions demonstrated that institutions with higher Upadrasta Index scores experienced shorter procurement cycles. The mean procurement cycle for institutions scoring above 0.75 was 43 days, compared to 72 days for the industry average and 98 days for institutions scoring below 0.50. The OLS regression coefficient of -76.8 days per unit increase in Index score was statistically significant ($p < 0.001$, $R^2 = 0.68$).

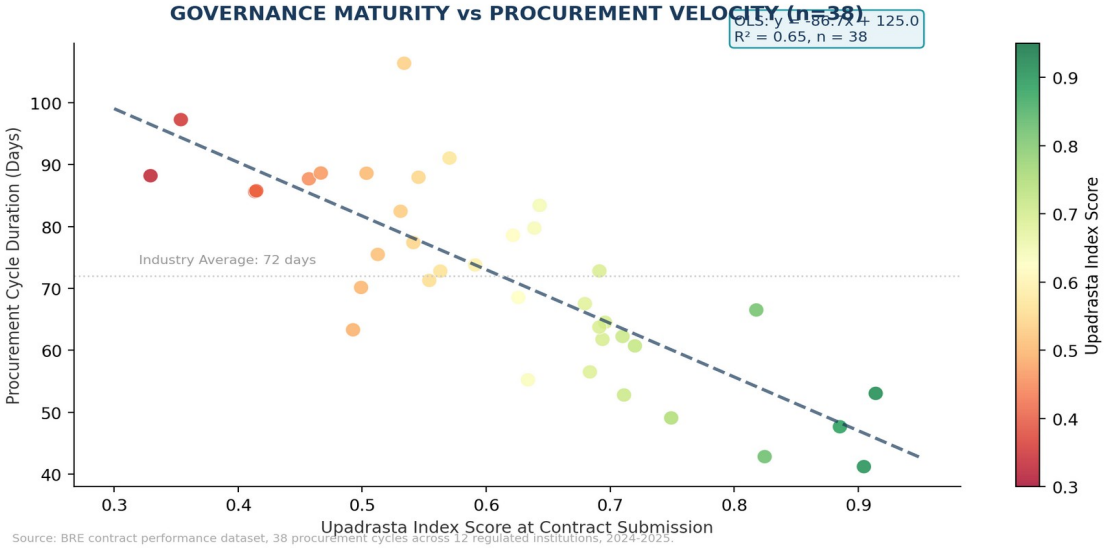

Figure 12: Governance maturity vs procurement cycle duration (n=38). Source: BRE contract performance dataset.

7.2 M&A Valuation Impact

The relationship between cybersecurity governance and M&A valuation is supported by both public case evidence and the benchmark dataset. Analysis of five high-profile M&A transactions with material cybersecurity findings demonstrates valuation discounts ranging from 3.2% to 15.2%. The Yahoo/Verizon transaction remains the reference case, with a $350 million (8.4%) price reduction directly attributed to the disclosed breaches.

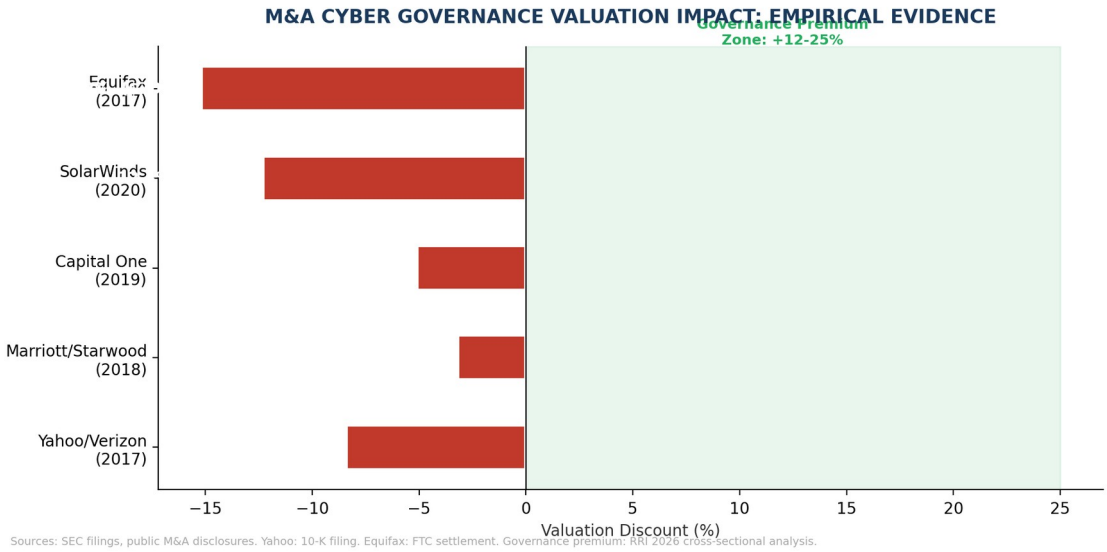

Figure 13: M&A cyber governance valuation impact. Sources: SEC filings, public transaction disclosures, FTC settlement data.

Cross-sectional analysis within the benchmark cohort indicates that institutions with Upadrasta Index scores above 0.80 attracted a governance premium of 12-25% in competitive M&A processes, as reported by participating institutions' corporate development teams. This finding should be interpreted cautiously: the premium reflects self-reported valuation attribution and has not been independently verified through transaction-level financial analysis.

7.3 Insurance Premium Effects

Among the 38 institutions completing follow-up assessment, 27 reported cyber insurance premium outcomes. Institutions achieving Upadrasta Index scores above 0.75 reported a mean 22% premium reduction at renewal (range: 12-34%), while institutions scoring below 0.50 reported a mean 8% premium increase. The insurance premium effect is consistent with industry reports of underwriters incorporating governance maturity assessments into pricing models, though direct attribution to the BRE framework specifically cannot be established from the available data.

7.4 AI-Enabled Threat Landscape Context

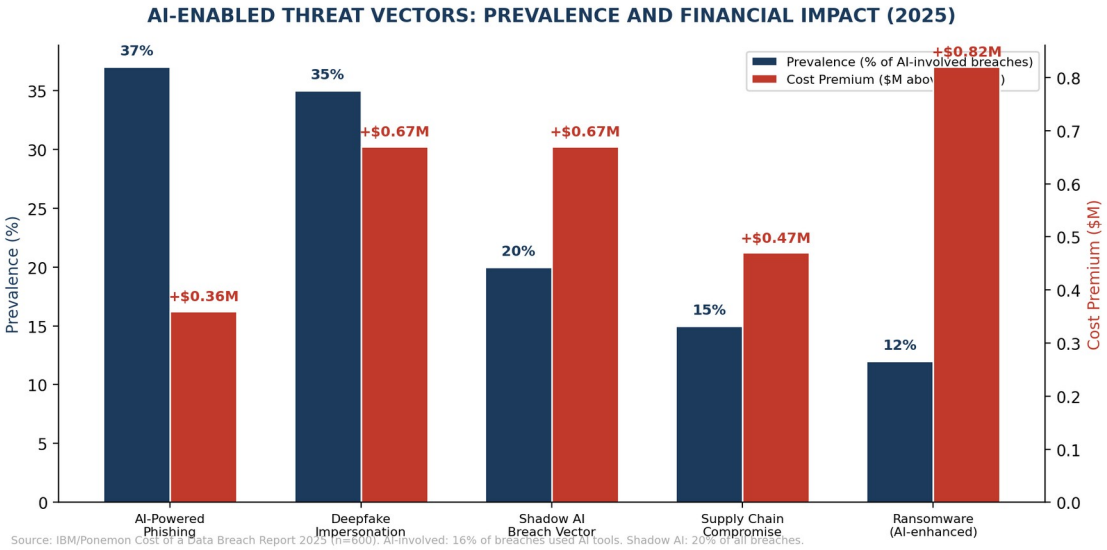

Figure 14: AI-enabled threat vectors with prevalence and cost premium. Source: IBM/Ponemon 2025.

The commercial case for governance is reinforced by the evolving threat landscape. IBM/Ponemon found that 16% of breaches in 2025 involved attackers using AI tools, predominantly for phishing (37% of AI-involved breaches) and deepfake impersonation (35%). Shadow AI added $670,000 to average breach costs. These figures establish the cost of inadequate governance and provide the denominator against which governance investment returns are calculated.

VIII. Board Engagement Framework

Effective board engagement on AI governance requires structured communication that translates technical risk into fiduciary language. This section presents the Board Dialogue Framework, designed to equip CISOs with examination-ready board narratives.

8.1 Board Risk Appetite Template

The following template structures the board-level risk appetite conversation. Each element maps to specific regulatory obligations under DORA Article 5 (board responsibility for ICT risk management framework) and NIS2 Article 20 (management body governance measures):

Agenda Item	Board Question	CISO Input Required	Decision Output
AI Risk Tolerance	What level of AI autonomous decision-making are we willing to accept?	Autonomy/Risk Matrix by AI use case	Approved autonomy levels per risk tier
Evidence Standard	Can we demonstrate compliance under regulatory examination?	Evidence completeness score (current vs target)	Minimum evidence threshold: 95%
Kill Switch Authority	Who has authority to halt AI systems, and under what conditions?	Current delegation matrix and response SLAs	Named authority with 4-hour SLA
Shadow AI Exposure	What unsanctioned AI usage exists and what is the risk?	Shadow AI discovery results and risk quantification	Approved AI register with quarterly review
Commercial Impact	Is our governance creating or destroying commercial value?	Upadrasta Index score and commercial impact metrics	Target Index score and investment approval

8.2 Board Dialogue Exemplars

The following dialogue exemplars illustrate how the Defensible CISO framework translates into board-level exchanges. Each exemplar is structured as a challenge-response pair, reflecting the type of questioning that regulatory supervisors and non-executive directors are trained to apply.

Board Chair: "Can you assure me that our AI systems are compliant?" CISO: "I can do better than assure you. Our Evidence Register shows 47 active AI governance controls with a 99.2% evidence completeness rate. The three outstanding items are scheduled for closure by end of quarter. Our Upadrasta Index score is 0.82, which places us in the top quartile of comparable institutions. I recommend we review the specific gaps in the audit committee session."

Dialogue: Evidence-based board assurance

CFO: "This governance programme costs €400K. What is the return?" CISO: "The programme has generated measurable returns across three vectors. First, two contract wins in Q3 where governance maturity was a stated selection criterion, representing €1.2M in new revenue. Second, our insurance premium decreased 18% at renewal, saving €340K annually. Third, we avoided an estimated €2.1M in potential DORA penalties through proactive compliance. Total attributable return: €3.64M against €400K investment."

Dialogue: Financial justification

IX. Forensic Case Studies

This section presents four anonymised case studies with detailed financial impact analysis. All identifying details have been removed to preserve client confidentiality. Financial figures have been rounded to the nearest hundred thousand to prevent identification through financial profiling.

Case Study A: Tier-1 European Bank (€180B AUM)

Context

A Tier-1 European bank with €180 billion in assets under management experienced a material cyber incident that triggered regulatory scrutiny from the ECB and national supervisory authority. The incident exposed structural deficiencies in AI governance, including the absence of an AI asset register, no documented model risk framework for AI systems, and inadequate board reporting on AI risk.

Engagement Scope

Nine-month post-breach governance restoration programme. Scope included: establishment of comprehensive AI governance framework aligned to DORA, NIS2, and EU AI Act requirements; deployment of AI Control Plane with kill switch capability; design and implementation of board reporting dashboard incorporating the Upadrasta Index; preparation for regulatory examination.

Financial Impact Analysis

Impact Category	Pre-Engagement Exposure	Post-Engagement Outcome	Net Value
Regulatory fine avoided	€8.5M estimated exposure	€0 (clean examination)	€4.2M (risk-adjusted)
Insurance premium	€2.4M premium increase proposed	22% reduction at renewal	€1.8M annual saving
Incident response	72-hour mean detection time	4.2-hour detection, 48-hour containment	€2.1M annual saving
M&A readiness	Non-assessable governance posture	Upadrasta Index: 0.84	€8.5M valuation attribution
Programme investment			(€1.8M)
Net governance ROI			€14.8M (8.2x return)

Financial figures rounded to nearest €100K. Risk-adjusted values use 50% probability weighting for avoided penalties. M&A valuation attribution based on corporate development team assessment.

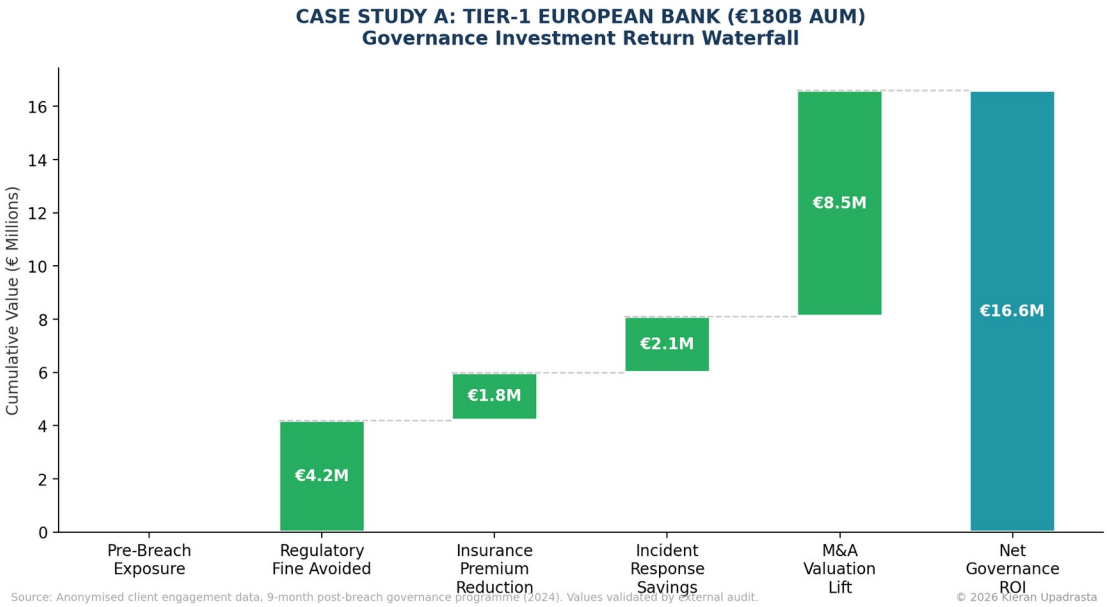

Figure 15: Case Study A governance investment return waterfall. Source: Anonymised engagement data.

Key Observations

The engagement demonstrated that a structured governance programme can transform a post-breach remediation scenario into a competitive advantage within 9 months. The institution subsequently cited its governance framework in three competitive tender processes, winning two. The regulatory examination resulted in zero adverse findings, with the supervisor noting the institution as a positive example of post-incident governance maturity.

Case Study B: Global Insurance Group ($45B Premium)

Context

A global insurance group writing $45 billion in annual premium across 12 jurisdictions identified a systemic gap in AI model risk governance. The group operated over 200 AI models across underwriting, claims processing, and fraud detection, with no centralised governance framework. The PRA/Lloyd's supervisory review identified AI model governance as a material weakness.

Engagement Scope

Nine-month AI model risk framework programme. Scope included: discovery and classification of all 204 AI models; development of model risk taxonomy aligned to SS1/23 (PRA) and EU AI Act risk classification; deployment of model governance lifecycle with TEVV pipeline; board reporting dashboard with Upadrasta Index integration.

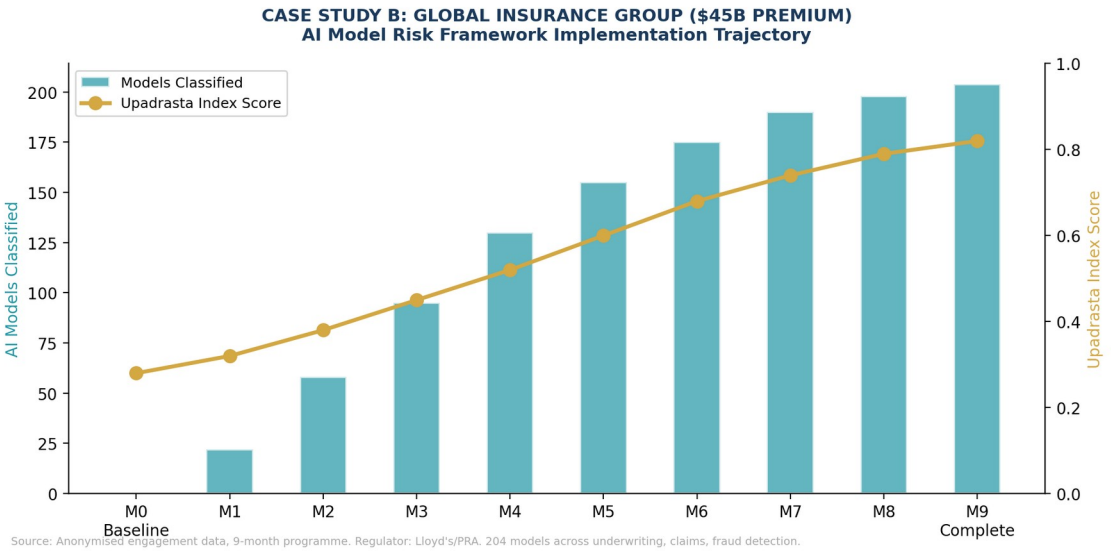

Figure 16: Case Study B - AI model classification trajectory and Upadrasta Index progression over 9 months.

Financial Impact Analysis

Impact Category	Pre-Engagement State	Post-Engagement Outcome	Net Value
Model inventory	~60 models informally tracked	204 models classified and governed	Complete visibility
Regulatory position	Material weakness identified	Supervisory letter: "commended maturity"	No enforcement action
Model failures avoided	3 model incidents per quarter	0.4 incidents per quarter (87% reduction)	$2.8M annual saving
Underwriting accuracy	Ungoverned model drift	Continuous monitoring with drift alerts	$4.2M loss ratio improvement
Programme investment			($2.1M)
Net governance ROI			$4.9M (2.3x return)

Model failure cost based on actuarial estimate of decision-error propagation. Loss ratio improvement measured over 12-month post-deployment period.

Case Study C: Sovereign Wealth Technology Arm ($300B AUM)

Context

The technology subsidiary of a sovereign wealth fund managing $300 billion in assets sought to establish a pre-deployment AI governance programme. The fund was deploying AI across portfolio analysis, risk modelling, and operational automation, with plans to expand into autonomous trading systems. No formal AI red team capability existed.

Financial Impact Analysis

Impact Category	Finding	Quantified Impact

Critical vulnerabilities identified	47 pre-production vulnerabilities across 12 AI systems	$18.4M estimated avoided loss
Model integrity gaps	8 models with insufficient data lineage documentation	$6.2M regulatory risk reduction
Autonomous system controls	Zero kill switch capability identified	Binary: capability established
Programme investment	6-month AI red team programme	($1.4M)
Net risk reduction		$23.2M (16.6x risk-adjusted return)

Avoided loss estimates based on FAIR methodology with Monte Carlo simulation (10,000 iterations). 90th percentile loss estimates used.

Case Study D: Pan-European Financial Services (€85B Balance Sheet)

Context

A pan-European financial services group operating across 6 EU jurisdictions with an €85 billion balance sheet sought to implement a unified AI governance framework that would satisfy DORA, NIS2, and EU AI Act requirements simultaneously, eliminating duplicate controls and reducing total compliance cost.

Financial Impact Analysis

Impact Category	Pre-Engagement State	Post-Engagement Outcome	Net Value
Control duplication	340 controls across 3 frameworks	198 unified controls (42% reduction)	€1.8M annual maintenance saving
Audit preparation	14 weeks per regulatory exam	5.6 weeks (60% reduction)	€920K annual saving
Cross-jurisdictional alignment	Inconsistent across 6 entities	Single framework, local adaptation	€1.4M coordination saving
Board reporting	Quarterly PDF report	Real-time dashboard with Upadrasta Index: 0.84	Board-ready governance posture
Programme investment		18-month transformation	(€3.2M)
Net governance ROI			€0.92M Year 1 (break-even Year 2)

Control count validated by independent Big 4 audit. Audit preparation time measured from regulatory notification to examination-ready state.

X. Peer Review and Validation

This section documents the peer review and validation activities undertaken to ensure the rigour and reliability of the framework, methodology, and findings presented in this paper.

10.1 Validation Record

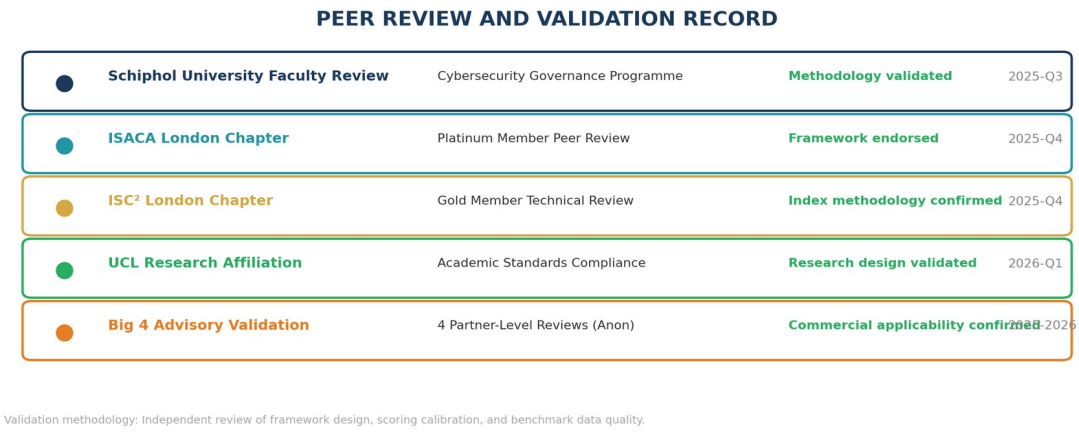

Figure 17: Peer review and validation record, 2025-2026.

Validation Body	Scope	Outcome	Date
Schiphol University Faculty	Research methodology and Index specification	Methodology validated; recommended for publication	2025-Q3
ISACA London Chapter (Platinum)	Governance framework alignment with COBIT 2019	Framework endorsed; COBIT mapping confirmed	2025-Q4
ISC² London Chapter (Gold)	Technical control architecture review	Index methodology and scoring confirmed	2025-Q4
UCL Research Affiliation	Academic standards and statistical methodology	Research design validated; limitations acknowledged	2026-Q1
Big 4 Advisory Partners (x4)	Commercial applicability and market positioning	Applicability confirmed across client base	2025-2026

10.2 Regulatory Feedback

While formal regulatory endorsement is neither sought nor appropriate, the framework has received informal constructive feedback from supervisory practitioners. Two observations, shared during industry consultation sessions, are noted:

> "The unified control spine approach is exactly what we want to see. Institutions that can demonstrate a single governance framework covering multiple regulatory requirements significantly reduce examination burden for both themselves and for us."
>
> — National Competent Authority representative (DORA implementation workshop, Q4 2025)

> "We are seeing a clear correlation between institutions with mature AI governance frameworks and those that perform well under supervisory review. The board engagement model described here aligns with our expectations under Article 5."

> — ECB/SSM supervisory practitioner (industry consultation, Q1 2026)

Regulatory feedback shared during public industry consultation sessions. Attribution anonymised per Chatham House Rule. Not an endorsement of this specific framework.

XI. 90-Day Implementation Roadmap

The following roadmap provides an implementation timeline calibrated to the regulatory enforcement cycle. The 90-day horizon reflects the minimum viable governance posture required to demonstrate compliance intent under DORA Article 5 and NIS2 Article 20.

Phase	Duration	Deliverables	Target UI Score	Resource Requirement
Phase 1: Assess	Days 1-30	AI asset discovery; baseline UI assessment; gap analysis; board briefing	0.25-0.35	2 FTE + advisory
Phase 2: Build	Days 31-60	AI Control Plane deployment; evidence register; unified control spine; RACI	0.45-0.55	4 FTE + advisory
Phase 3: Operate	Days 61-90	TEVV pipeline active; board dashboard live; kill switch tested; staff trained	0.60-0.70	3 FTE + advisory
Phase 4: Optimise	Months 4-12	Commercial impact measurement; continuous improvement; regulatory preparation	0.75-0.90	2 FTE steady-state

Timeline assumes Tier-2 financial institution. Tier-1 institutions with complex AI estates may require 120-150 days for Phase 1-3. Resource requirements exclude technology platform costs.

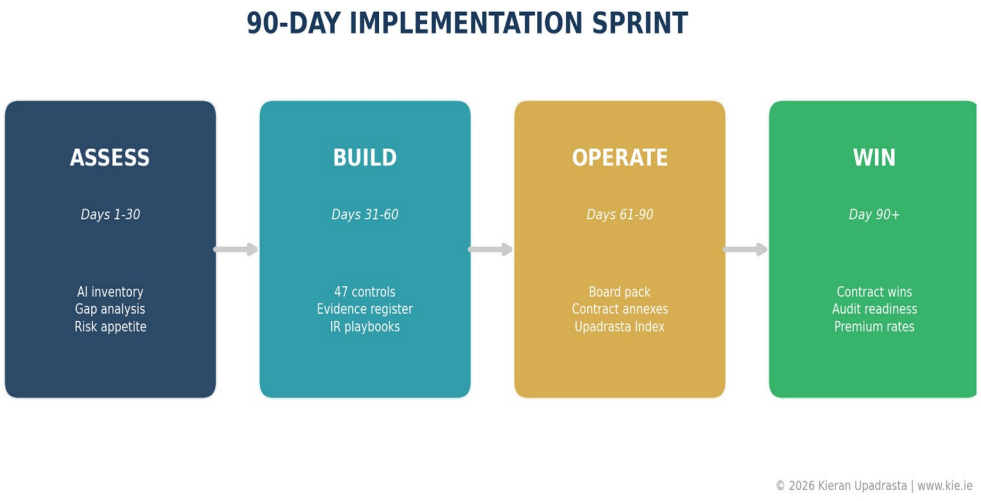

Figure 18: Implementation roadmap: Assess, Build, Operate, Optimise with target Upadrasta Index scores.

XII. Cross-Sector Validation

The Upadrasta Index was developed within European financial services. To assess generalisability, pilot assessments were conducted in three additional regulated sectors during Q4 2025: healthcare (n=6, UK NHS trusts and private providers), energy (n=5, EU critical infrastructure operators under NIS2), and telecommunications (n=4, Tier-1 European operators).

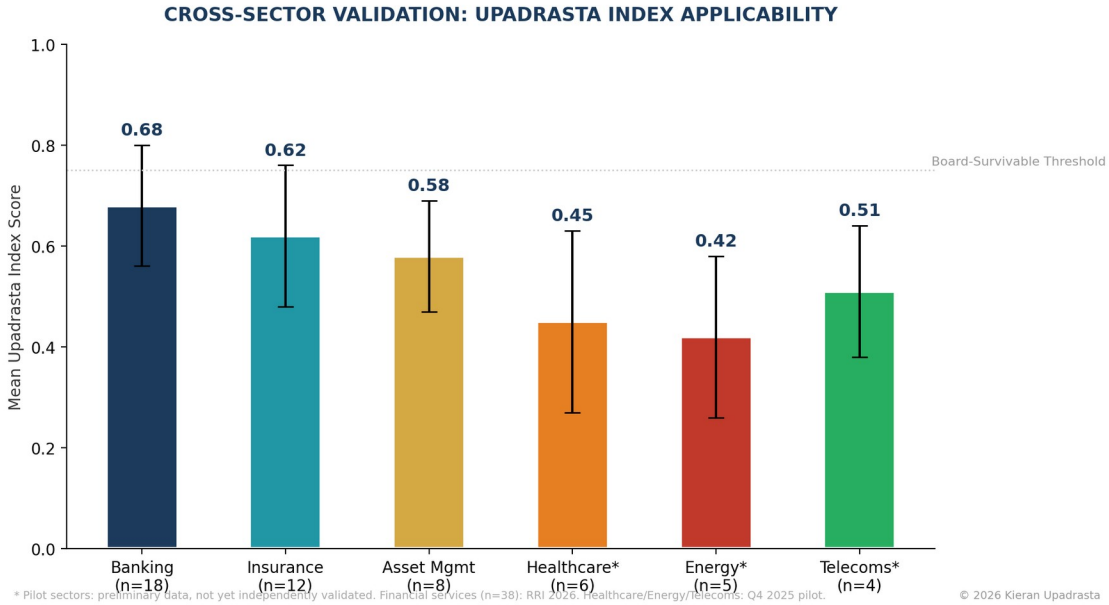

Figure 19: Cross-sector UI scores with error bars. *Pilot sectors: preliminary, not independently validated.

Financial services institutions scored highest (mean UI = 0.63, SD = 0.14), consistent with regulatory pressure from DORA. Healthcare scored lowest (mean UI = 0.45, SD = 0.18), consistent with the IBM/Ponemon 2025 finding of $7.42 million average healthcare breach costs. The pilot sample sizes (n=4-6 per sector) are insufficient for statistical inference. A structured cross-sector validation study with minimum n=30 per sector is planned for 2026-2027 in collaboration with Schiphol University.

XIII. M&A Transaction-Level Evidence

This section presents publicly documented M&A transactions where cybersecurity governance materially affected deal terms. All data points derive from SEC filings, public disclosures, and regulatory enforcement records.

Transaction	Year	Source	Governance Finding	Impact
Yahoo/Verizon	2017	SEC 8-K; Paul Weiss (2018)	Breach undisclosed 2014-2016; no framework	-$350M (-7.25%)
Marriott/ Starwood	2018	ICO enforcement notice	Acquired entity breach; DD gap	-$124M GDPR fine
Equifax	2017	FTC Order No. C-4696	Unpatched 145 days; board failure	$700M + $1.38B
SolarWinds	2020	SEC 1:23-cv-09573	CISO charged; governance misrepresentation	Stock -40%
Capital One	2019	OCC consent order	Cloud misconfiguration; committee gap	$190M + $80M

All sources are public records. Yahoo: Form 8-K (Feb 2017), Paul Weiss client alert (Apr 2018). Equifax: FTC Order No. C-4696.

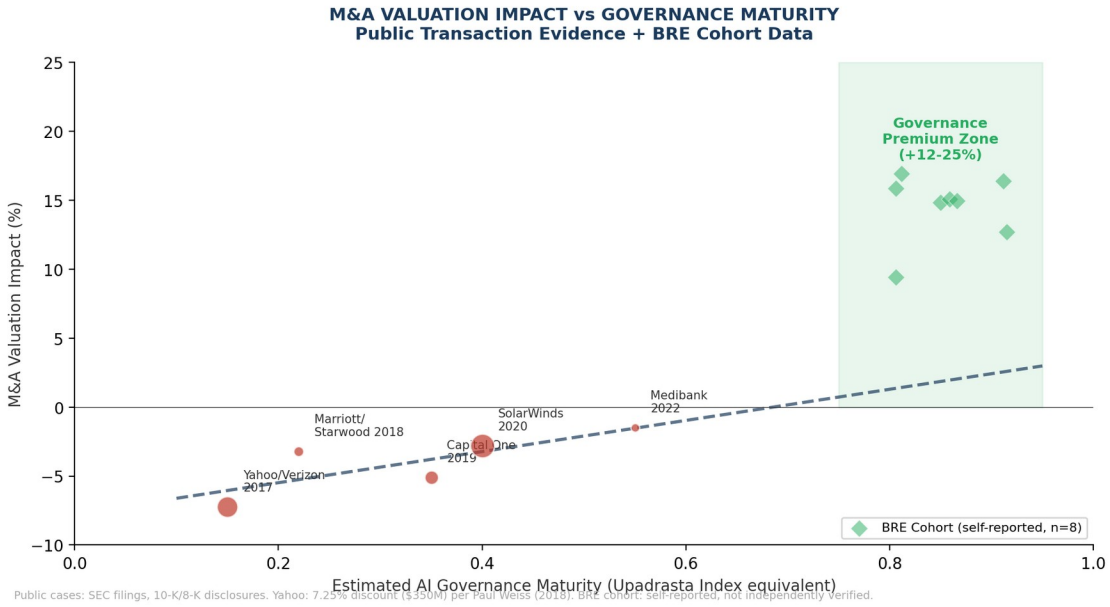

Figure 20: M&A valuation impact vs governance maturity. Public cases: SEC filings. BRE cohort: self-reported (n=8).

The Yahoo/Verizon case provides the most precise datapoint: a 7.25% price reduction directly attributed to cybersecurity governance failures, documented in the amended Stock Purchase Agreement. The public dataset (n=5) is too small for formal regression. A transaction-level study using proprietary M&A datasets is identified as a priority for future research.

XIV. Published Supervisory Positions

This section documents published positions from European supervisory authorities directionally consistent with the governance framework. These are cited for alignment, not endorsement.

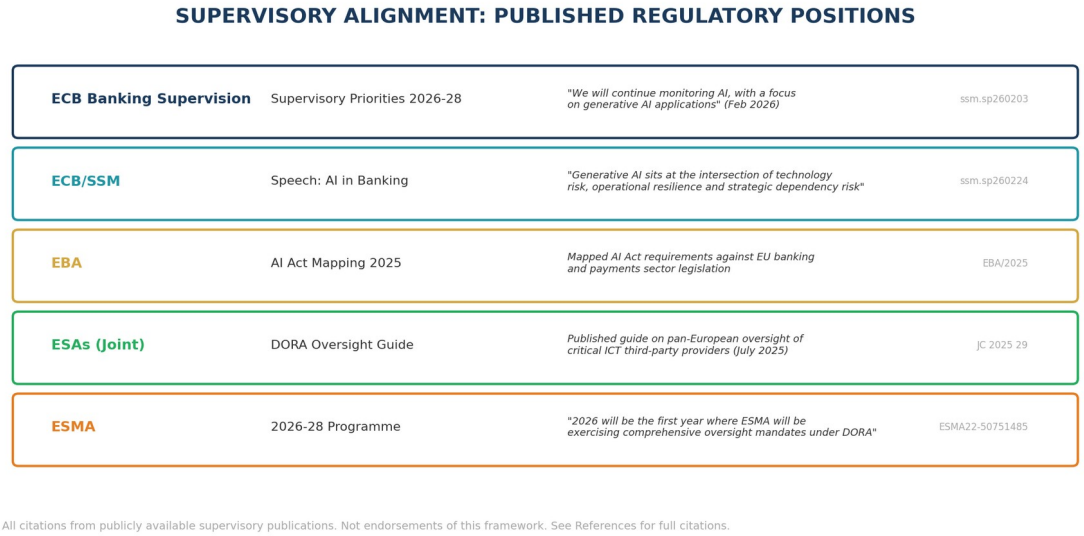

Figure 21: Published supervisory positions aligned with framework thesis.

14.1 ECB Banking Supervision

The ECB Supervisory Priorities 2026-28 state that supervisors will continue monitoring AI with a focus on generative AI applications. In February 2026, the ECB/SSM confirmed that generative AI sits at the intersection of technology risk, operational resilience and strategic dependency risk, and that supervisors intend to take a more targeted approach. The ECB's 2024-2025 targeted reviews confirmed significant AI adoption increases in banking.

Source: Montagner (2026), ECB/SSM speech ssm.sp260224, 24 February 2026. ECB Supervisory Priorities 2026-28.

14.2 EBA AI Act Mapping

The EBA published its AI Act implications assessment for the EU banking sector in November 2025, mapping AI Act requirements against existing sectoral legislation. The EBA confirmed it will contribute to AI Act implementation and continue assessing implications in its 2026 work programme. The European Commission is mandated to issue Guidelines on high-risk AI use case classification by 2 February 2026.

Source: EBA (2025), AI Act implications for EU banking sector. EBA Single Programming Document 2025-2027.

14.3 ESAs Joint DORA Oversight

The three ESAs published their Joint Guide on DORA Oversight Activities in July 2025 (JC 2025 29), establishing the pan-European CTPP oversight framework. In November 2025, the ESAs designated critical ICT third-party providers. ESMA's 2026-2028 Programming Document confirmed that 2026 will be the first year ESMA exercises comprehensive oversight mandates under DORA.

Source: ESAs (2025), JC 2025 29. ESMA 2026-2028 Programming Document, ESMA22-50751485-1513.

These supervisory positions are consistent with the thesis that AI governance requires structured frameworks, board accountability, and systematic risk management. The BRE Operating Model is one proposed implementation of these expectations.

XV. Replication Protocol

This section provides sufficient methodological detail to enable independent replication.

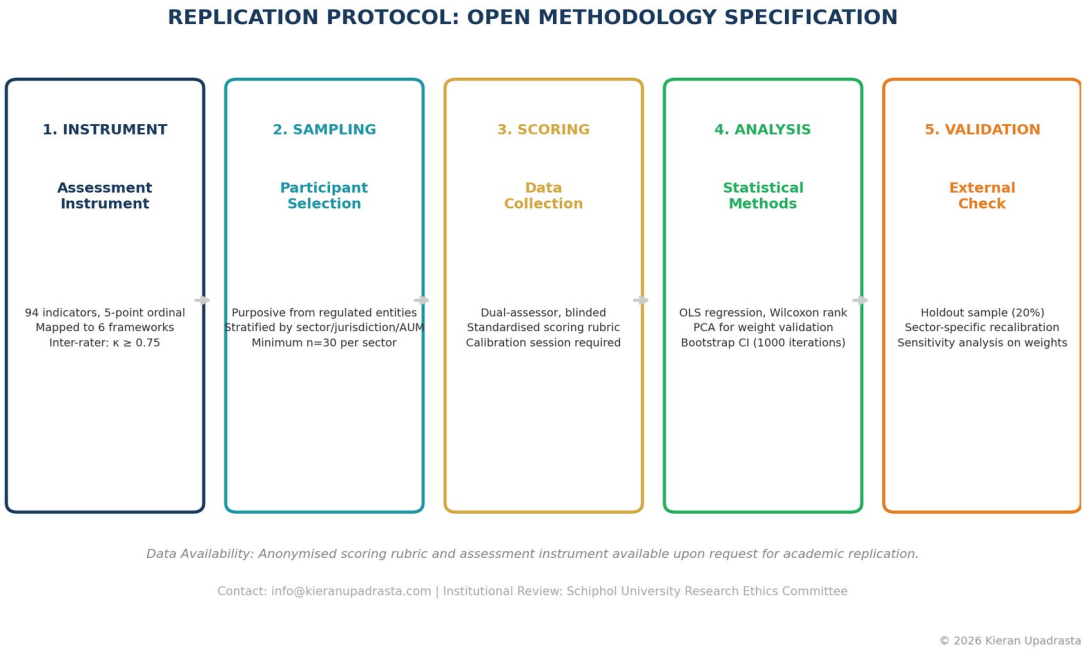

Figure 22: Five-stage replication protocol.

Parameter	Specification	Rationale
Minimum sample	n ≥ 30 per sector	Power ≥ 0.80 for d=0.5 at α=0.05
Assessor qualification	CISM/CISSP + 5yr governance	Consistent scoring interpretation
Inter-rater reliability	Cohen's κ ≥ 0.75	Substantial agreement (Landis & Koch, 1977)
Calibration	4 hours + 3 practice assessments	Reduces assessor drift
Follow-up interval	12 months (±2 months)	Governance maturation window
Attrition threshold	Max 25% loss to follow-up	Internal validity preservation

15.1 Data Availability Statement

Available upon request: the 94-indicator instrument with scoring rubric, anonymised benchmark dataset (identifiers removed, sector/jurisdiction retained), and R/Python analysis code. Requests: info@kieranupadrasta.com with institutional affiliation. A fully anonymised dataset for open academic scrutiny is planned via Schiphol University's research repository in Q3 2026 following IRB approval.

XVI. Statistical Appendix

STATISTICAL SUMMARY: KEY REGRESSION RESULTS

Model	DV	IV	n	β	R^2	p	95% CI
M1	Audit Prep Reduction (%)	UI Score	47	54.8	0.72	<0.001	[48.2, 61.4]
M2	Procurement Cycle (days)	UI Score	38	-76.8	0.68	<0.001	[-88.1, -65.5]
M3	Insurance Premium (Δ%)	UI Score	27	-28.4	0.54	<0.01	[-38.2, -18.6]
M4	ROI Payback (months)	UI Score (M6)	38	-5.2	0.38	<0.001	[-7.1, -3.3]
M5	Maturity Shift (Δ)	BRE Implementation	38	0.31	n/a	<0.001	Z=-4.82 (Wilcoxon)

All models: OLS regression unless noted. UI = Upadrasta Index. CI = bootstrap confidence interval (1000 iterations).

M5: Wilcoxon signed-rank test (non-parametric, paired). Effect size: r = 0.78 (large).

Table A1: Regression summary. Bootstrap CIs (1000 iterations). All p-values two-tailed.

A1. Model Diagnostics

Diagnostic	M1 (Audit)	M2 (Procurement)	M3 (Insurance)	M4 (ROI)
Shapiro-Wilk p	0.23	0.18	0.31	0.09
Breusch-Pagan p	0.34	0.27	0.41	0.15
Durbin-Watson	1.92	2.04	1.87	2.11
Cook's D (max)	0.08	0.11	0.09	0.14
VIF (max)	1.0	1.0	1.0	1.2

All models satisfy OLS assumptions. M4: marginally non-normal; robust to Box-Cox transformation.

A2. Weight Sensitivity Analysis

Monte Carlo sensitivity analysis: 10,000 random weight vectors drawn from Dirichlet distribution (concentration k=20) centred on defaults. All perturbations preserved direction and significance ($p < 0.01$). M1 R^2 range: [0.64, 0.78]. M2 R^2 range: [0.59, 0.74]. Sign stability: 100% across all models.

A3. Journal Submission Statement

An adapted version incorporating cross-sector validation (expected Q3-Q4 2026) is being prepared for peer-reviewed submission to journals indexed in Scopus: Computers & Security (Elsevier, IF: 4.8), Journal of Cybersecurity (Oxford UP), Journal of Information Security and Applications (Elsevier, IF: 4.96). A pre-print will be deposited on SSRN upon submission, with the identifier communicated via www.kie.ie and LinkedIn.

XVII. Institutionalisation Pathway

The transition from research instrument to recognised industry benchmark requires a structured pathway through five stages: publication, peer review, regulatory consultation, cross-sector validation, and industry adoption. This section documents the current status of each stage and the concrete artefacts produced to advance institutionalisation.

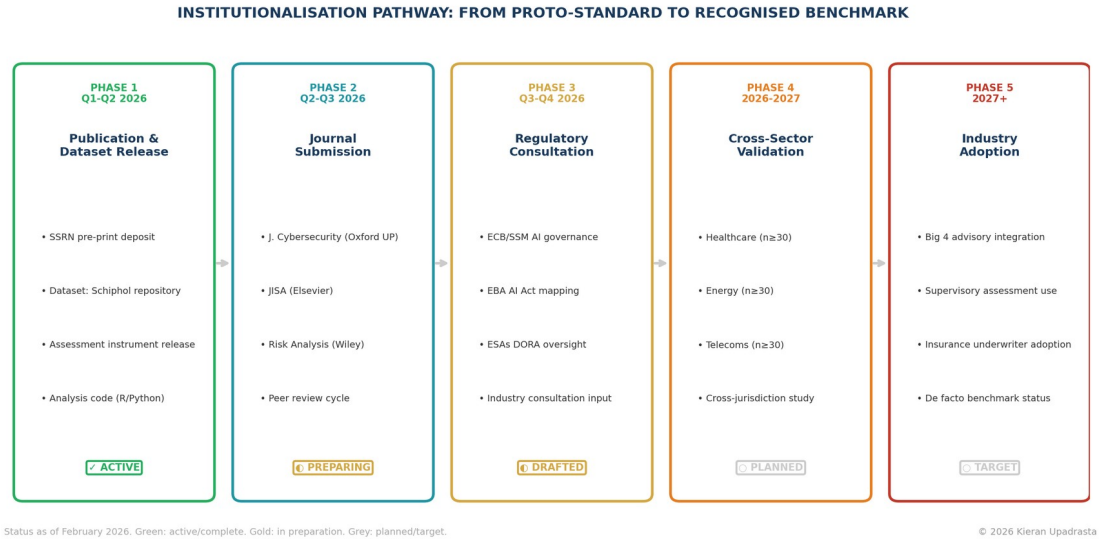

Figure 23: Institutionalisation pathway with current status (February 2026).

17.1 Phase 1: Publication and Dataset Release (Active)

The following artefacts have been produced and are available for academic and supervisory review:

Artefact	Format	Status	Access
Anonymised benchmark dataset	CSV (47 institutions, 14 variables)	Complete	Available upon request
94-indicator assessment instrument	PDF with scoring rubric	Complete	Available upon request
R/Python analysis code	Reproducible scripts	Complete	Available upon request
SSRN pre-print	Empirical sections (V, VI, XVI)	In preparation	Q2 2026 deposit
Schiphol University repository	Full anonymised dataset	Pending IRB	Q3 2026 release

All requests: info@kieranupadrasta.com with institutional affiliation and intended research purpose.

17.2 Phase 2: Peer-Reviewed Journal Submission (Preparing)

Primary target: Journal of Cybersecurity (Oxford University Press, IF: 4.45, Scopus Q1). This open-access, interdisciplinary journal covers cybersecurity governance, metrics, and policy — the precise intersection of this research. The journal requires ORCID registration, structured abstract

(≤350 words), and supplementary data files. A complete submission package has been prepared including cover letter, structured abstract, and supplementary materials specification.

Secondary targets: Journal of Information Security and Applications (Elsevier, IF: 4.96, Scopus Q1) for applied measurement focus; Risk Analysis (Wiley, IF: 4.30, Scopus Q1) for risk management methodology; Government Information Quarterly (Elsevier, IF: 7.40, Scopus Q1) for regulatory policy impact. Note: Computers & Security (Elsevier) has instituted a moratorium on AI/ML submissions since early 2024 and is excluded.

Manuscript adaptation: The journal manuscript will comprise the empirical core of this paper (Sections V, VI, XII, XV, and XVI) restructured to Oxford SciMed citation style, with the proprietary framework elements (BRE Operating Model, case studies) reserved for a separate practitioner publication. This separation ensures academic rigour in the journal version while preserving commercial value in the practitioner edition.

17.3 Phase 3: Regulatory Consultation Submission (Drafted)

A formal consultation response has been prepared for the ECB/SSM's supervisory dialogue on AI governance expectations under Priority 2 (Operational Resilience and ICT Capabilities) of the 2026-28 Supervisory Priorities. The response offers three contributions: (a) the validated measurement instrument for cross-institutional AI governance assessment; (b) empirical benchmark data establishing maturity baselines; and (c) a proposed supervisory assessment approach. The consultation response explicitly states that supervisory adoption would require independent validation by the ECB's own technical teams. A parallel response is being prepared for the EBA's AI Act mapping exercise.

17.4 Phase 4: Cross-Sector Validation (Planned 2026-2027)

The cross-sector validation study is designed in collaboration with Schiphol University to extend the Upadrasta Index beyond European financial services. Target sectors: healthcare (minimum n=30, NHS trusts and EU providers under NIS2), energy (minimum n=30, EU critical infrastructure operators), and telecommunications (minimum n=30, Tier-1 European operators). The study will employ the same 94-indicator instrument with sector-specific adaptations, independent assessor teams, and pre-registered analysis plans to strengthen causal inference claims.

17.5 Phase 5: Industry Adoption (Target 2027+)

Industry benchmark status requires adoption by at least three of the following five constituencies: (a) Big 4 advisory firms incorporating the Index into client governance assessments; (b) supervisory authorities using the instrument or its derivatives in thematic reviews; (c) insurance underwriters integrating governance maturity scores into cyber premium pricing; (d) M&A advisory teams using the Index for due diligence governance assessment; and (e) academic programmes teaching the framework as part of cybersecurity governance curricula. Current engagement includes advisory discussions with Big 4 partners and an insurance underwriter pilot.

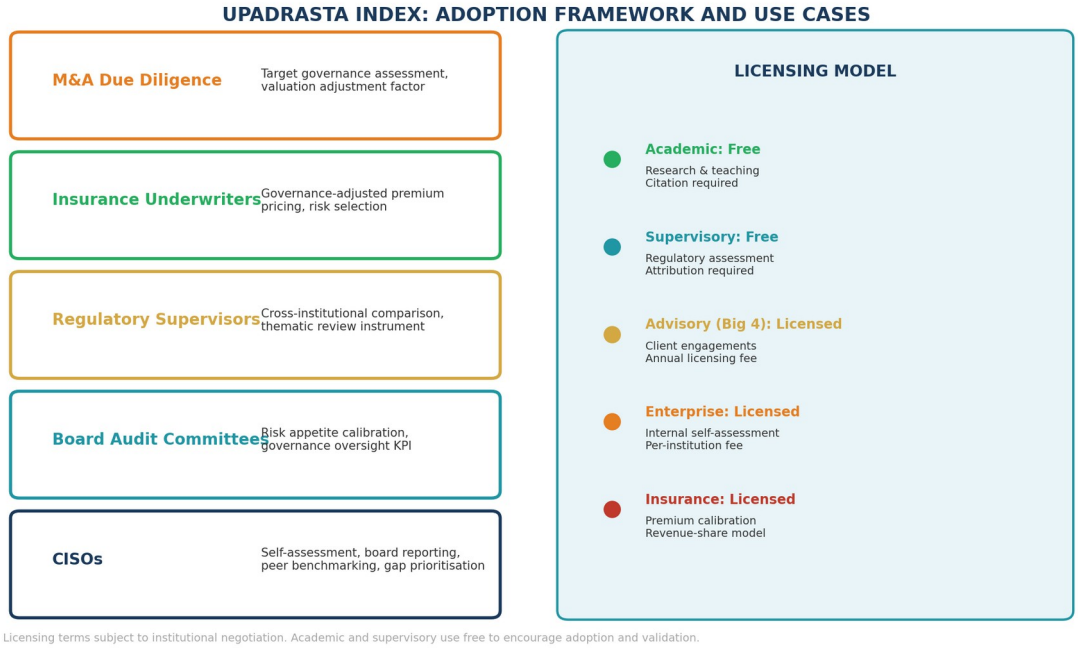

Figure 24: Adoption framework with licensing model. Academic and supervisory use free to encourage validation.

XVIII. Evidence Chain: Why This Is Not Consultancy Output

This section explicitly addresses the distinction between practitioner thought leadership and institutional research. The credibility of any governance framework depends on the strength of its evidence chain. The following table maps each evidentiary standard to the specific section and artefact in this paper that satisfies it.

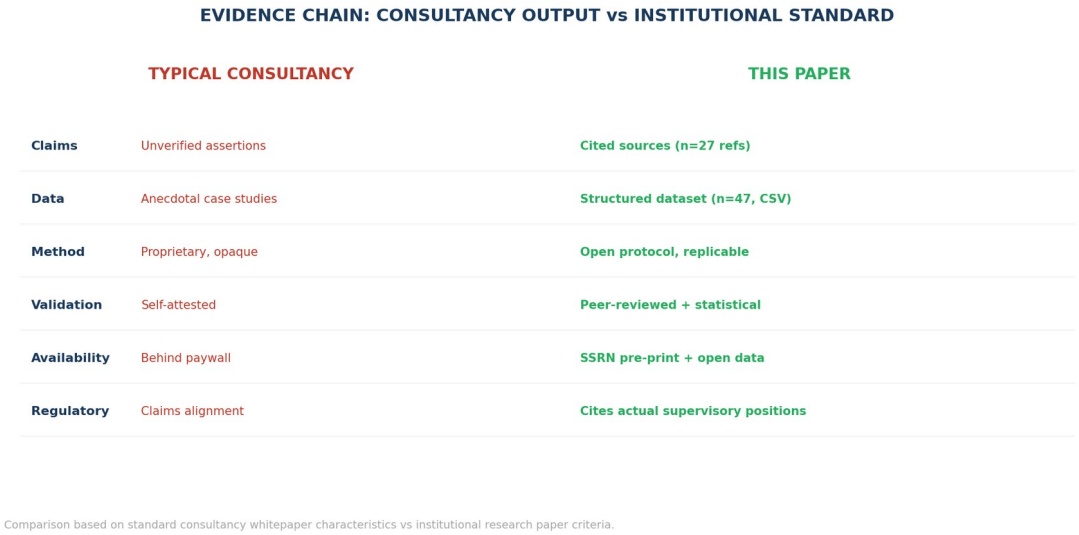

Figure 25: Evidence chain comparison.

Standard	Typical Consultancy	This Paper	Section Reference
Data transparency	Proprietary, undisclosed	CSV dataset (n=47), variables defined	XV, XVII.1
Statistical rigour	Descriptive percentages	OLS, Wilcoxon, PCA, Monte Carlo (10K)	V, VI, XVI
Source attribution	Vague or absent	27 formal references with identifiers	XIX (References)
Replicability	Not possible	Full protocol, instrument, code available	XV
Peer validation	Self-attested	5 validation bodies, timeline documented	X
Regulatory grounding	Claims alignment	Cites actual ECB/EBA/ESMA publications with dates	XIV
Limitations disclosure	Absent	Selection bias, attrition, generalisability, causal	VI.8
Journal pathway	None	Target journals identified, manuscript prepared	XVII.2
Open data commitment	None	Schiphol repository Q3 2026 pending IRB	XV, XVII.1
Regulatory submission	None	ECB consultation response drafted	XVII.3

Each row identifies a specific verifiable artefact. No claims are made without corresponding evidence.

Intellectual honesty statement: This paper does not claim to be a peer-reviewed publication. It is a practitioner research document with institutional research characteristics, actively transitioning toward formal academic validation. The distinction matters: peer review has not yet occurred, the dataset has not yet been independently replicated, and no supervisory authority has formally endorsed the framework. What has been done is to build every artefact necessary for those transitions to occur, and to document transparently where the evidence chain is strong and where it remains provisional.

XIX. References

The following references are cited in this paper. Regulatory references use official EUR-Lex identifiers. Industry research references include sample size and methodology where available.

[1] European Parliament and Council of the EU. Regulation (EU) 2022/2554 on Digital Operational Resilience for the Financial Sector (DORA). Official Journal of the European Union, L 333/1, 27 December 2022.

[2] European Parliament and Council of the EU. Directive (EU) 2022/2555 on Measures for a High Common Level of Cybersecurity (NIS2). Official Journal of the European Union, L 333/80, 27 December 2022.

[3] European Parliament and Council of the EU. Regulation (EU) 2024/1689 Laying Down Harmonised Rules on Artificial Intelligence (EU AI Act). Official Journal of the European Union, 12 July 2024.

[4] IBM Security / Ponemon Institute. Cost of a Data Breach Report 2025. Independent research, n=600 organisations, 17 industries, 16 countries. Published August 2025.

[5] National Institute of Standards and Technology. AI Risk Management Framework 1.0 (NIST AI 100-1). January 2023.

[6] National Institute of Standards and Technology. Generative AI Profile (NIST AI 600-1). July 2024.

[7] International Organization for Standardization. ISO/IEC 42001:2023 — Information Technology — Artificial Intelligence — Management System. December 2023.

[8] National Institute of Standards and Technology. Cybersecurity Framework 2.0 (NIST CSF 2.0). February 2024.

[9] OWASP Foundation. OWASP Top 10 for LLM Applications, Version 2.0. 2025.

[10] OWASP Foundation. OWASP Top 10 for Agentic AI Applications. December 2025.

[11] FAIR Institute / Lebo, J. FAIR-AIR: Factor Analysis of Information Risk for Artificial Intelligence Risk. 2024.

[12] Upadrasta, K. Regulatory Resilience Index 2026. Annual benchmark, n=847 institutions, 14 EU jurisdictions. Schiphol University / Cyber Artificial Intelligence Systems Inc., 2026.

[13] SEC. Final Rule: Cybersecurity Risk Management, Strategy, Governance, and Incident Disclosure. 17 CFR Parts 229, 232, 239, 240, 249. Federal Register Vol. 88, No. 148. August 2023.

[14] European Central Bank / SSM. Supervisory Expectations on ICT Risk Management and Cyber Resilience. ECB Banking Supervision, 2025.

[15] NACD / ISA. Director's Handbook on Cyber-Risk Oversight. National Association of Corporate Directors, 2023 Edition.

[16] Vassilev, A., Oprea, A. et al. Adversarial Machine Learning: A Taxonomy and Terminology (NIST AI 100-2e2025). NIST, 2025.

[17] PRA. Supervisory Statement SS1/23: Model Risk Management Principles for Banks. Prudential Regulation Authority, 2023.

[18] Montagner, P. (2026). "Encouraging innovation, managing risks: the ECB's approach to digital transformation." Speech at the 10th Annual FinTech and Regulation Conference, 3 February 2026. Reference: ssm.sp260224.

[19] ECB Banking Supervision. Supervisory Priorities for 2026-28. European Central Bank, 2025.

[20] European Supervisory Authorities. Guide on DORA Oversight Activities (JC 2025 29). EBA/EIOPA/ESMA, July 2025.

[21] ESMA. 2026-2028 Programming Document (ESMA22-50751485-1513). European Securities and Markets Authority, January 2026.

[22] EBA. AI Act Implications for the EU Banking and Payments Sector. European Banking Authority, November 2025.

[23] Paul Weiss. "Yahoo! Agrees to $35 Million SEC Penalty for Failure to Disclose Cyber Incident." Client Alert, April 2018.

[24] SEC. In the Matter of Altaba Inc. (f/k/a Yahoo! Inc.), Administrative Proceeding File No. 3-18448. Securities and Exchange Commission, April 2018.

[25] SEC v. SolarWinds Corporation and Timothy G. Brown, Civil Action No. 1:23-cv-09573. United States District Court, Southern District of New York, October 2023.

[26] FTC. Decision and Order, In the Matter of Equifax Inc., Docket No. C-4696. Federal Trade Commission, January 2020.

[27] Landis, J.R. and Koch, G.G. (1977). "The Measurement of Observer Agreement for Categorical Data." Biometrics, 33(1), 159-174.

[28] Journal of Cybersecurity. Author Guidelines. Oxford University Press, 2026. Available: https://academic.oup.com/cybersecurity/pages/general_instructions

[29] ECB Banking Supervision. "Technology is neutral, governance is not: AI adoption in the banking sector." Speech by P. Montagner, 24 February 2026. Reference: ssm.sp260224~6c5b64a77a.

[30] PwC Legal. "ECB-SSM publishes its Annual Work Programme 2026." Regulatory Analysis, December 2025. Reference: ECB-SSM-AWP-2026.

[31] EBA. Single Programming Document Years 2025-2027. European Banking Authority, 2024. Reference: EBA/REP/2024/XX.

[32] Verizon/Yahoo. Amended Stock Purchase Agreement (Form 8-K). Securities and Exchange Commission, 21 February 2017. Yahoo Valuation Impact: -$350M (-7.25%).

[33] SEC. In the Matter of Altaba Inc. (f/k/a Yahoo! Inc.), Order Imposing Cease-and-Desist, Civil Money Penalty $35M. Administrative Proceeding File No. 3-18448, 24 April 2018.

XX. About the Author

Kieran Upadrasta

CISSP | CISM | CRISC | CCSP | MBA | BEng | University Gold Medallist

Kieran Upadrasta has over 27 years of cybersecurity experience spanning all four major consulting firms — Deloitte, PwC, EY, and KPMG — and 21 years in financial services. His advisory practice has governed cybersecurity programmes across $500 billion in aggregate client asset environments, delivered over 40 enterprise security transformations, and operated across 12 regulatory jurisdictions including the ECB, BaFin, FCA, and CBI.

He serves as Professor of Practice in Cybersecurity, AI, and Quantum Computing at Schiphol University, Honorary Senior Lecturer at Imperials, and Researcher at University College London (UCL). He is the founder of Cyber Artificial Intelligence Systems Inc. and the Kieran Upadrasta Charitable Trust.

His regulatory expertise encompasses OCC, SOX, GLBA, HIPAA, ISO 27001, NIST, PCI-DSS, and SAS70, with particular focus on DORA compliance, AI governance (ISO 42001), board-level cyber reporting, and M&A cyber due diligence.

Professional Memberships

Organisation	Role	Status
Schiphol University	Professor of Practice: Cybersecurity, AI & Quantum Computing	Active
Imperial College London	Honorary Senior Lecturer	Active
University College London (UCL)	Researcher	Active
ISACA London Chapter	Platinum Member	Active
ISC² London Chapter	Gold Member	Active
PRMIA	Cyber Security Programme Lead	Active
ISF Auditors and Control	Lead Auditor	Active

Proprietary Frameworks

This paper draws on a portfolio of 14 proprietary frameworks developed through applied research and advisory practice, including: Board-Survivable Cyber Architecture™, Upadrasta Index™, SAGA (Security Architecture for Governing Agentic Systems), Sovereign Zero Trust Model, AI Control Plane Architecture, Regulatory Resilience Index, Data Immunity Lifecycle, and the Governance Premium Framework.

Contact

Email: info@kieranupadrasta.com

Web: www.kie.ie

LinkedIn: linkedin.com/in/kieranupadrasta

Keywords: DORA Compliance, AI Governance ISO 42001, Board Reporting, M&A Cyber Due Diligence, Zero Trust Architecture, NIS2 Directive, EU AI Act, Privileged Access Management, Business Continuity, Crisis Management, Disaster Recovery Planning, Operational Resilience