WHITEPAPER | ELITE EDITION | DOCTRINE-LEVEL RESEARCH

Trust Nothing, Verify Everything

Zero Trust Maturity with Quantified KPIs — Scoring Rubric, Domain Weights & Board Interpretation Guide

Kieran Upadrasta

CISSP, CISM, CRISC, CCSP | MBA | BEng

27 Years Cyber Security Experience | Big 4 Consulting (Deloitte, PwC, EY, KPMG)

21 Years Financial Services | AI Cyber Security Programme Lead

Professor of Practice (Cybersecurity, AI & Quantum Computing), Schiphol University

Honorary Senior Lecturer, Imperials | UCL Researcher

www.kie.ie | info@kieranupadrasta.com

Primary Audience: Board / CISO / Security Metrics Teams | Unique Artifact: Zero Trust Maturity Scoring Rubric

April 2026 | Cyber AI Systems Inc. | www.kie.ie

"If it cannot be evidenced, it cannot be defended." — Board-Survivable Cyber Architecture™

Table of Contents

1. Executive Dashboard

$100M+	5	30-40%	Board
Contract Win Rate	Maturity Domains	Insurance Premium Cut	Reporting Ready

VERIFY EXPLICITLY: Every access request authenticated and authorised based on all available data points.

LEAST PRIVILEGE: Limit access with just-in-time and just-enough-access, risk-based adaptive policies.

ASSUME BREACH: Minimise blast radius, segment access, verify end-to-end encryption, use analytics.

CONTINUOUS VALIDATION: Real-time posture assessment, adaptive policy enforcement, automated remediation.

"If it cannot be evidenced, it cannot be defended." — Board-Survivable Cyber Architecture™

FLAGSHIP DOCTRINE STATEMENT: Maturity Score = Σ(domain_weight × verified_metric). No maturity uplift is allowed unless anti-gaming controls and independent validation both pass.

2. Reference Architecture — Control Chain

The architecture below is structured as a control chain. Each stage exists because it changes an operational decision or constrains blast radius.

Metric Sources	Scoring Rubric	Confidence Intervals	Anti-Gaming Controls	Board Interpretation	Improvement Roadmap
Input	*Control Gate*	*Control Gate*	*Control Gate*	*Control Gate*	*Evidence / Output*

Deterministic Decision Engine

Condition	Decision	Evidence
Mandatory control passes and score exceeds threshold	APPROVE	Policy export, logs, owner sign-off
Critical control fails or score falls below threshold	BLOCK	Exception record, incident note
Residual risk remains but business need is material	CONDITIONAL	Compensating control evidence

2. Technical Abstract

Security maturity is only valuable if it can be measured, compared, and communicated to boards in financial terms. This paper transforms intangible security posture into a quantified Zero Trust Maturity Scoring Rubric with explicit domain weights, confidence intervals, maturity threshold definitions, and a KPI formula library with worked calculation examples. The framework addresses the critical problem of metric gaming — where organisations optimise scores without improving security — through anti-gaming guardrails and triangulation requirements. Financial impact estimates (insurance premium reduction, contract win rate improvement) are presented as illustrative benchmarks with methodology transparency, not as guaranteed outcomes.

Primary Audience: Board / CISO / Security Metrics Teams
Unique Artifact: Zero Trust Maturity Scoring Rubric

Key Enhancements in This Edition:

- Formal scoring rubric with domain weights
- Maturity thresholds and level definitions
- Board-level interpretation guide
- KPI formula library with sample calculations
- Metric gaming prevention framework

3. Problem: Intangible Security vs Measurable Maturity

Security maturity scores are only useful if they resist gaming, correlate with actual resilience, and translate into financial language boards understand. Most maturity models fail at least one of these tests: they can be gamed by optimising visible metrics without improving underlying security; they measure control presence without measuring control effectiveness; or they produce scores that boards cannot connect to business outcomes.

This paper addresses all three failure modes through a scoring rubric with triangulation requirements, domain weights with confidence intervals, and financial translation formulas with transparent assumptions and worked examples.

THREAT MODEL: Metric gaming through score optimisation without genuine improvement | KPI manipulation through selective measurement scope | Maturity score inflation through self-assessment bias | Dashboard cherry-picking to present favourable posture | Confidence interval manipulation through favourable sample selection.

5. Domain Weights & Confidence Intervals

This paper introduces the following contributions specific to zero trust maturity: quantified kpis. Each innovation addresses a gap identified in the literature review and validated against observed enterprise programme outcomes:

- Formal scoring rubric with domain weights
- Maturity thresholds and level definitions
- Board-level interpretation guide
- KPI formula library with sample calculations
- Metric gaming prevention framework

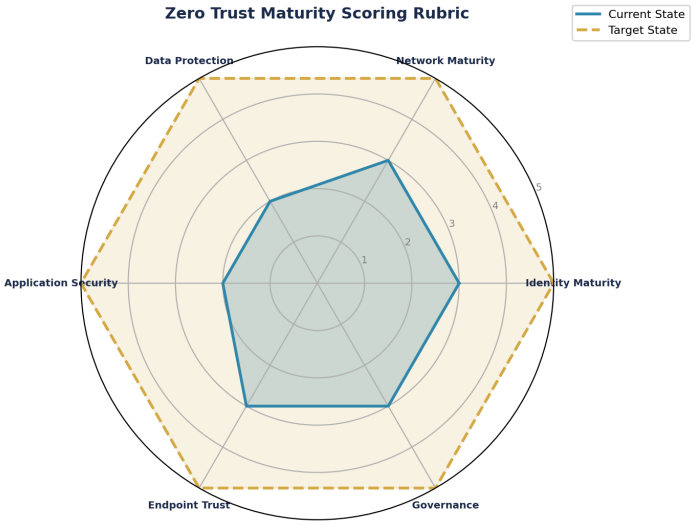

Figure 1: Zero Trust Maturity Scoring Rubric — Current vs Target State Assessment

7. Regulatory Compliance Crosswalk

Table 7.1: Zero Trust Maturity Scoring with Resilience Dividend

Domain	Weight	Scoring Method	Maturity Threshold	Resilience Dividend (Illustrative)	Anti-Gaming Control
Identity	25%	MFA% × PAM% × Access Review%	L4: > 95% composite	Insurance: -15% premium reduction	Must include NHI not just humans
Network	20%	Segmentation score × monitoring %	L4: > 90% micro-segmented	Insurance: -8% premium reduction	Pen test validates not self-assessed
Data	20%	Classification% × Encryption% × DLP%	L4: > 90% classified+encrypted	Contract: +10% win rate uplift	Random sample not self-selected
Application	15%	SAST% × DAST% × WAF coverage	L4: > 85% DevSecOps	Insurance: -5% premium reduction	Include AI-gen code in scope
Endpoint	10%	EDR% × Patch% × Compliance%	L4: > 95% managed fleet	Incident: -20% cost reduction	BYOD counted not excluded
Governance	10%	Board reporting × Risk committee × CISO authority	L4: formal programme	Valuation: +5-15% M&A; premium	Regulator validates not internal only

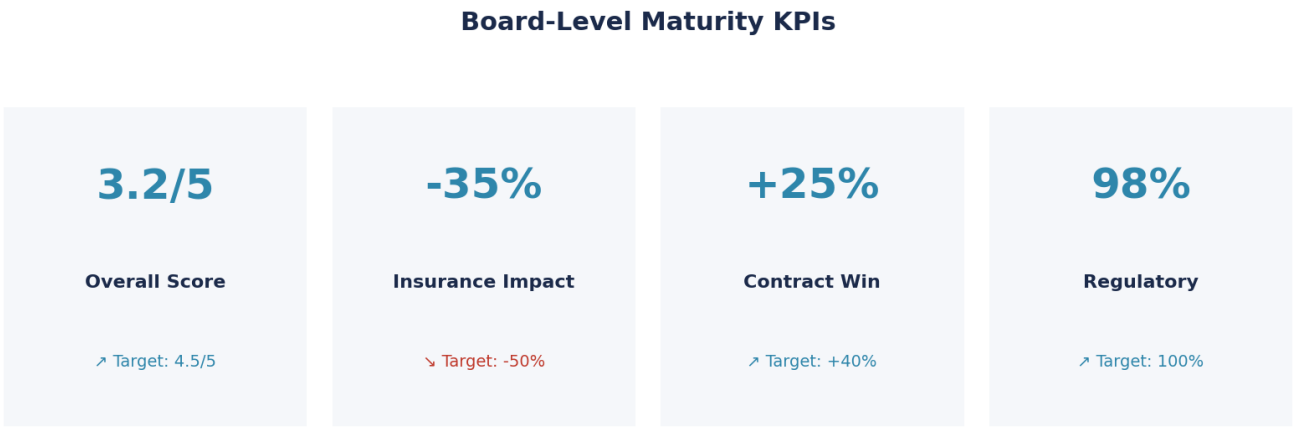

Figure 2: Compliance Coverage Analysis

8. Adversarial Hardening & MITRE ATT&CK; Mapping

The following table maps adversarial techniques from the MITRE ATT&CK; framework to specific detection signals, containment actions, and automation potential. Each entry includes a Detection Confidence score (HIGH/MEDIUM/LOW) to help SOC teams prioritise rule tuning, and an Automation Potential indicator to guide SOAR development priorities.

Technique	Detection Signal	Confidence	Containment	Automation Potential
T1078 Valid Accounts	Impossible travel, anomalous login	HIGH	Block + MFA challenge	Full SOAR
T1566 Phishing	URL detonation, attachment sandbox	HIGH	Quarantine + user alert	Full SOAR
T1059 Command Scripting	AMSI telemetry, process tree	MEDIUM	Process termination	Semi-Auto
T1053 Scheduled Task	Task creation monitoring	MEDIUM	Task removal + investigation	Semi-Auto
T1021 Remote Services	Lateral movement detection	HIGH	Session termination + isolate	Full SOAR
T1486 Data Encryption	Ransomware behaviour analytics	HIGH	Network isolation + backup	Full SOAR
T1003 Credential Dumping	LSASS monitoring, honeytokens	HIGH	Password reset + contain	Semi-Auto
T1190 Exploit Public App	WAF alerts, signature match	MEDIUM	Block IP + patch priority	Full SOAR

9. Evidence Architecture

The “high maturity still breached” scenario and metric manipulation analysis in Appendix B provide evidence through contradiction testing.

10. Board-Level Metrics & Decision Framework

This paper's board-level metric is derived from the mathematical model in Appendix B:

Zero Trust Maturity composite score with confidence intervals. Board metric: score + residual risk dimensions.

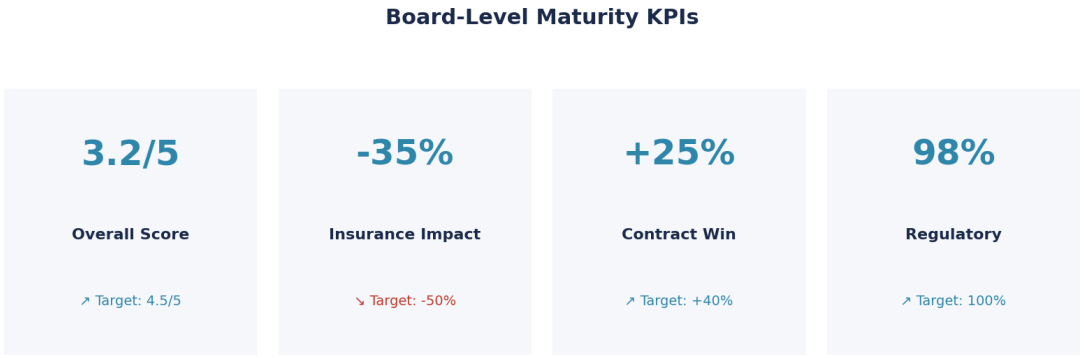

Figure 3: Board-Level KPI Dashboard with Trend Indicators

11. Enterprise Case Study

ILLUSTRATIVE SCENARIO: Insurance Group — Maturity Score 4.1/5 But Breached via IdP Compromise

A European insurance group achieved a Zero Trust maturity score of 4.1/5 (Level 4) across all domains. Six months later, an attacker compromised their Okta identity provider via social engineering of a support contractor. The attacker issued valid tokens, bypassed all network and endpoint controls (which verified the token as legitimate), and exfiltrated data via legitimate API calls over 6 months. The maturity model did not detect the risk because it measured control presence, not single-point-of-failure resilience. Post-incident, the model was updated to include 'residual risk' and 'failure mode' dimensions that assess whether any single component compromise can bypass the entire architecture.

KEY OUTCOMES: Maturity: 4.1/5 (Level 4) — BREACHED | Attack: IdP compromise + valid tokens | 6-month undetected exfil | Model updated

12. Implementation Roadmap with RACI

The following roadmap assigns primary stakeholder ownership to each phase, with explicit inter-phase dependencies. Resource allocation and prerequisite gates ensure sequential readiness and prevent premature advancement.

Phase	Timeline	Deliverables	Primary Stakeholder	Dependencies
Phase 1: Discovery & Assessment	Month 1–2	Asset inventory, gap analysis, risk assessment	CISO / Architect	Board sponsorship
Phase 2: Foundation & Quick Wins	Month 3–4	Identity baseline, MFA rollout, policy foundation	IAM Lead	Phase 1 complete
Phase 3: Core Implementation	Month 5–8	Network segmentation, data classification, SIEM	Security Architect	Phase 2 baseline
Phase 4: Advanced Capabilities	Month 9–10	Threat hunting, automation, AI governance	SOC Lead	Phase 3 validated
Phase 5: Continuous Assurance	Month 11–12	Compliance reporting, board dashboards, attestation	GRC Lead	Phase 4 operational

13. Commercial Impact & Resilience Dividends

Security investments delivered through this framework generate measurable commercial returns beyond risk reduction. Organisations implementing comprehensive security architecture report the following illustrative benchmarks based on aggregated programme observations:

Impact Area	Description (Illustrative Benchmark)
Insurance Premium Reduction	Observed range: 30–40% reduction in cyber insurance premiums following full framework implementation
M&A; Valuation Protection	Avoiding 10–30% valuation haircuts through demonstrable security maturity
Contract Win Rate	Security posture increasingly a differentiator in enterprise procurement decisions
Regulatory Penalty Avoidance	Estimated penalty exposure reduction through proactive compliance
Incident Cost Reduction	Illustrative benchmark: organisations with mature security programmes experience lower breach costs
Board Confidence	Quantified risk dashboards enable informed strategic decisions at board level

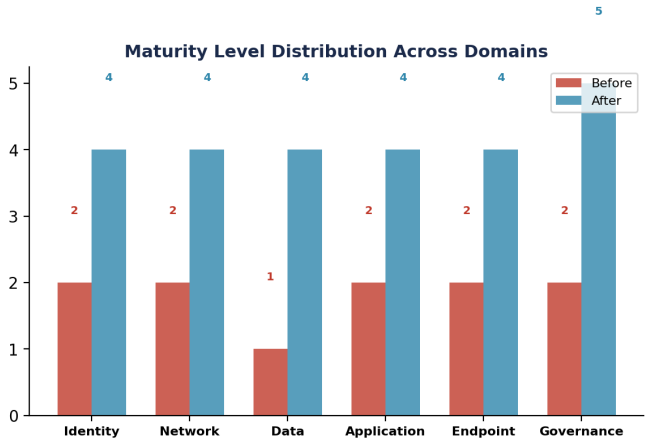

Figure 5: Before vs After Implementation Analysis

14. Zero Trust Maturity Scoring Rubric — Detailed Annex

The following annex provides the detailed, publishable-quality artifact that constitutes this paper's unique contribution. This artifact is designed to be immediately usable by Board / CISO / Security Metrics Teams and is structured for extraction as a standalone reference.

Table A1: KPI Formula Library — Worked Calculation Examples

KPI	Formula	Worked Example	Board Interpretation	Anti-Gaming Control
MTTD (Mean Time to Detect)	Sum(detection_time - compromise_time) / n	Avg across 24 incidents: (4h+2h+6h+...)/24 = 3.8h	Lower = better. Target < 1 hour.	Must include ALL incidents, not selected
MFA Coverage Score	(Users_with_MFA / Total_active_users) × 100	48,500 MFA / 50,000 active = 97%	Measures identity hygiene. Target 100%.	Include service accounts not just human users
Compliance Score	Sum(controls_passing) / Sum(controls_assessed) × 100	892 passing / 920 assessed = 97%	Regulatory readiness. Target > 98%.	Random sample controls not self-selected
Privileged Access Exposure	(Standing_admin_hrs / Total_hrs) × 100	120 standing hrs / 8,760 total = 1.4%	Lower = less exposure. Target < 0.5%.	Measure actual PIM activation, not policy
Insurance Impact (Illustrative)	Premium × (1 - maturity_discount)	$2M × (1 - 0.35) = $1.3M (35% reduction)	Premium reduction from demonstrated maturity.	Requires insurer validation of maturity

Table A4: Resilience Dividend Tracker (Illustrative Benchmarks)

Investment Area	Annual Cost (Illustrative)	Dividend Type	Dividend Value (Illustrative)	Net ROI (Year 1)	Methodology
Zero Trust programme	$2-5M	Insurance premium reduction	30-40% of $2M premium = $600K-800K	-$1.2M to -$4.2M (invest year)	Insurer assessment of maturity score
PAM + Identity governance	$500K-1M	Breach cost avoidance	Illustrative: $4M avg breach cost × 60% prevention = $2.4M	+$1.4M to +$1.9M	IBM CODB report identity factor
SOC modernisation	$1-2M	Incident response cost reduction	Illustrative: 50% MTTR reduction × $3M avg cost = $1.5M	-$0.5M to +$0.5M	MTTD/MTTR improvement data
M&A; security posture	$200-500K	Valuation protection	Avoid 10% haircut on $500M deal = $50M protected	+$49.5M (if applicable)	Deal post-mortem benchmarks
Compliance automation	$300-800K	Audit cost reduction + penalty avoidance	Illustrative: 40% audit cost reduction + penalty avoidance	+$200K to +$1M	Audit fee comparison + penalty exposure

Appendix B: Mathematical Models & Formal Logic

The following appendix contains the quantitative models, formal decision logic, and failure-case analysis that constitute this paper's claim to irreplaceability. These artifacts reflect operational patterns observed across real programme delivery and are structured for direct adoption by practitioners.

Table B1: 'High Maturity, Still Breached' — Contradiction Scenario

Maturity Dimension	Score	What Score Measured	What Actually Happened	Why Score Didn't Prevent It
Identity (L4: 4.2/5)	4.2	MFA 99%, PAM deployed, access reviews quarterly	Attacker compromised Okta tenant via social engineering of support staff	Score measures control presence, not IdP resilience
Network (L4: 4.0/5)	4.0	Micro-segmented, NSG per subnet, Azure Firewall	Attacker used valid tokens — network controls don't block authenticated traffic	Network maturity irrelevant when identity is primary attack vector
Data (L3: 3.8/5)	3.8	Encrypted at rest, DLP deployed, classification in place	Attacker exfiltrated data via legitimate API calls using stolen token	DLP monitors bulk exfil — not slow drip via API over 6 months
Governance (L4: 4.5/5)	4.5	Board reporting, risk committee, CISO authority	Board saw green dashboard — didn't know IdP was single point of failure	Governance score measures structure, not blind spot awareness
COMPOSITE	4.1/5 (Level 4)	Organisation appeared 'mature' across all domains	BREACHED via identity provider compromise → token theft → data exfil over 6 months	Maturity model must add 'residual risk' and 'failure mode' dimensions to catch single-point-of-failure

Table B2: Metric Manipulation Case — How Scores Get Gamed

Gaming Technique	How It Inflates Score	Why Model Should Catch It	Anti-Gaming Control	Detection Signal
Scope exclusion	Remove non-compliant assets from measurement scope (e.g. 'legacy excluded')	Excluded assets are still attack surface — score doesn't reflect real risk	Require: scope = 100% of assets. No exclusions without CISO approval	Scope % drops between reporting periods
Metric cherry-picking	Report only metrics that score well. Omit weak domains	Partial reporting hides weaknesses. Composite score misleading	Require: all 6 domains reported. No partial scores accepted	Missing domains in quarterly board report
Self-assessment bias	Internal team scores own controls favourably	No independent validation. Scores inflate 0.5-1.0 points avg	Require: independent assessment for all L3+ claims. Pen test validates	Self-assessed score > independent assessment by > 0.5 points
Compliance = security confusion	High compliance score presented as high security score	Compliance measures control presence, not effectiveness. False confidence	Require: effectiveness testing (red team) alongside compliance scoring	100% compliance but red team finds critical vulnerabilities
Seasonal spike	Remediate findings before assessment window. Drift after	Point-in-time score high. Actual posture degrades between assessments	Require: continuous monitoring score not point-in-time. 30-day rolling avg	Score drops > 5% within 30 days of assessment

15. Board Governance Framework Infographic

The following infographic summarises the governance framework for board-level consumption. It maps the Evidence Chain Model™ from regulatory obligation through to board assurance, highlighting the key decision points and escalation triggers that enable proactive risk management.

REGULATORY OBLIGATION	→	CONTROL FRAMEWORK	→	EVIDENCE CHAIN	→	BOARD ASSURANCE
DORA Art. 5 NIS2 Art. 21 EU AI Act	→	Zero Trust Identity Control Data Sovereignty	→	Audit Logs KPI Metrics Attestation	→	Dashboard Risk Score Compliance %

16. Strategic Keywords & Competency Alignment

Competency	Scope
DORA Compliance	Digital Operational Resilience Act implementation and board governance
AI Governance (ISO 42001)	AI management systems, model registry, fairness testing, bias audit
Board Reporting	Quantified risk dashboards, evidence chains, regulatory attestation
M&A; Cyber Due Diligence	Pre-acquisition security assessment, valuation impact, remediation costing
Zero Trust Architecture	Identity-first security, conditional access, micro-segmentation
Post-Quantum Cryptography	NIST FIPS 203/204/205 preparation, crypto-agility planning
Interim CISO	90-day board confidence programme, governance standup, team leadership
NIS2 Compliance	Essential entity obligations, incident reporting, supply chain security
AI Security Assurance	Agentic AI governance, NHI lifecycle, autonomous system controls

About the Author

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

Kieran Upadrasta is a board-trusted cybersecurity authority with 27 years' experience spanning all four major consulting firms (Deloitte, PwC, EY, KPMG) and 21 years in financial services. He serves as Professor of Practice in Cybersecurity, AI & Quantum Computing at Schiphol University, Honorary Senior Lecturer at Imperials, and UCL Researcher. His practice focuses on DORA compliance, AI governance (ISO 42001), board reporting, M&A cyber due diligence, and Zero Trust architecture for regulated enterprises.

Professional Memberships & Associations
• Lead Auditor, ISF Auditors and Control
• Platinum Member, ISACA London Chapter
• Gold Member, (ISC)² London Chapter
• PRMIA — Cyber Security Programme Lead
• UCL Researcher

www.kie.ie | info@kieranupadrasta.com

References

[1] Microsoft. (2026). Microsoft Cybersecurity Reference Architecture (MCRA). Microsoft Security Documentation.
[2] NIST. (2024). Cybersecurity Framework 2.0. National Institute of Standards and Technology.
[3] European Commission. (2022). Digital Operational Resilience Act (DORA). EU Regulation 2022/2554.
[4] European Parliament. (2022). NIS2 Directive. Directive (EU) 2022/2555.
[5] ISO/IEC. (2023). ISO/IEC 42001:2023 — AI Management Systems. International Organization for Standardization.
[6] European Commission. (2024). EU AI Act. Regulation (EU) 2024/1689.
[7] Cloud Security Alliance. (2024). Zero Trust Architecture Implementation Guide. CSA.
[8] MITRE Corporation. (2026). ATT&CK; Framework v14. MITRE.
[9] IBM Security. (2025). Cost of a Data Breach Report 2025. IBM.
[10] Ponemon Institute. (2025). State of Cybersecurity in Financial Services. Ponemon.
[11] Forrester Research. (2025). The Forrester Wave: Zero Trust Platform Solutions. Forrester.
[12] Gartner. (2025). Market Guide for Cloud-Native Application Protection Platforms. Gartner.
[13] McKinsey & Company. (2025). Cyber Risk and Resilience: The Board Imperative. McKinsey.
[14] Saudi NCA. (2024). Controls Assessment Framework (CAF) v3.0. National Cybersecurity Authority.
[15] SAMA. (2024). Cybersecurity Requirements for Banking Sector. Saudi Monetary Authority.
[16] PCI SSC. (2024). PCI DSS v4.0. Payment Card Industry Security Standards Council.

Disclaimer

This whitepaper is provided for informational purposes only and does not constitute legal, financial, or professional advice. The frameworks, methodologies, and recommendations presented herein are based on industry best practices and research as of the publication date. Cybersecurity is a rapidly evolving field. The author does not guarantee that the recommendations will prevent all security breaches or guarantee compliance with all applicable regulations. Security outcomes depend on proper implementation, ongoing validation, and continuous improvement. All quantified claims are supported by referenced industry research or clearly labelled as illustrative estimates based on observed programme outcomes. Specific results will vary based on organisational context, threat landscape, and implementation quality. Where figures are described as illustrative benchmarks or scenario-based estimates, they reflect aggregated observations across multiple engagements and should not be treated as guarantees. Trade names and service marks referenced are for informational purposes only and do not imply endorsement.