TOP 1% CYBERSECURITY WHITEPAPER

The AI-Driven Threat Frontier:

Zero Trust, Identity & Supply Chain Resilience

A Security Leader's Roadmap for 2026 and Beyond

Kieran Upadrasta

CISSP, CISM, CRISC, CCSP | MBA | BEng

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

21 Years Financial Services | Honorary Senior Lecturer

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

DORA Compliance | AI Governance (ISO 42001) | Zero Trust | Board Reporting | M&A Cyber Due Diligence

Table of Contents

Board Cyber Governance Pack

This one-page summary provides the board-ready framework for cyber governance oversight. Use quarterly for strategic review.

KPI	Target	Red Threshold	Owner	Cadence
Mean Time to Detect (MTTD)	<24 hours	>72 hours	CISO	Monthly
Mean Time to Respond (MTTR)	<4 hours	>24 hours	CISO	Monthly
Tier-0 Asset MFA Coverage	>98%	<90%	IAM Lead	Monthly
NHI Standing Privilege Ratio	<5%	>20%	IAM Lead	Quarterly
Critical Patch Currency (<7 days)	>95%	<80%	Infra Lead	Monthly
Third-Party IR Playbook Coverage	>80%	<50%	TPRM Lead	Quarterly
Phishing Resilience Rate	>90%	<70%	Security Ops	Quarterly
Regulatory Compliance Score	100%	<95%	GRC Lead	Quarterly

Board Reporting Cadence

- Quarterly: Full cyber risk report with KPI dashboard, regulatory status, incident summary
- Annually: Independent assessment, penetration test results, cyber insurance review
- Event-driven: Material incidents within 24 hours, regulatory findings within 48 hours

Five Questions for Every Board Meeting

1. What is our current cyber risk exposure in dollar terms? (FAIR methodology)
2. Are we detecting and responding faster than adversary breakout time? (51 seconds benchmark)
3. What percentage of our machine identities have standing admin privileges?
4. Have our critical third parties tested incident response with us this quarter?
5. What regulatory deadlines are approaching and what is our compliance gap?

Method & Evidence Standard

This section documents the research methodology and evidence standards applied throughout this whitepaper, enabling readers to assess the reliability and applicability of each claim.

Source Classification Taxonomy

Tier	Source Type	Treatment	Examples
Tier 1	Regulatory text, peer-reviewed research, official statistics	Cited directly with article/page reference	DORA regulation, SEC filings, NIST frameworks
Tier 2	Independent research firms with disclosed methodology	Cited with publication date, sample size noted	Verizon DBIR, IBM X-Force, Mandiant M-Trends
Tier 3	Vendor-sponsored research, surveys with potential bias	Cross-validated against Tier 1-2; limitations noted	Entro NHI Report, Silverfort surveys

Handling Conflicting Statistics

When sources report different figures for the same phenomenon:

- Report the range with attribution: 'Supply chain attacks affected 54-75% of organizations (Synopsys 54%, BlackBerry 75%)'
- Prioritize methodology transparency: Sources disclosing sample size, collection method, and date range are preferred
- Distinguish incidence from perception: The paper corrects the '85% deepfake' statistic (threat perception) vs. 49-50% (actual incidents)
- Flag projections vs. measured data: The $60B supply chain cost is labeled as Cybersecurity Ventures projection, not measured cost

Corrections Applied in This Paper

Commonly Cited	Corrected Version	Primary Source
144:1 NHI ratio (CyberArk)	144:1 ratio from Entro Labs	Entro H1 2025 Report, July 22, 2025
$350M Yahoo regulatory fine	$35M SEC fine; $350M was acquisition price reduction	SEC Order, April 24, 2018
85% deepfake incidents	85% perception; 49-50% actual incidents	Regula Deepfake Trends 2024
$60B supply chain cost	Projection, not measured cost	Cybersecurity Ventures, 2024

Executive Summary

FOUR DEFENSIBLE CLAIMS

1. Attack velocity now requires sub-minute detection: CrowdStrike documents 51-second breakout (February 2025).

2. Non-human identities are the unmanaged attack surface: 144:1 machine-to-human ratio with 97% having excessive privileges (Entro, July 2025).

3. Third-party risk doubled in one year: 15% to 30% of breaches now involve third parties (Verizon DBIR 2025).

4. AI-enabled attacks are operational reality: First documented AI-orchestrated campaign reported November 2025 (Anthropic).

On November 13, 2025, Anthropic published what security researchers characterized as the first publicly documented AI-orchestrated cyber espionage campaign, attributed to Chinese state-sponsored group GTG-1002. This campaign—active since at least September 2025—demonstrated that threat actors could deploy AI to perform 80-90% of attack operations autonomously.

This whitepaper synthesizes findings from 27 years of Big 4 consulting experience across financial services, combined with primary research from authoritative sources. Unlike derivative analyses, this document provides traceability (every major statistic cited to primary source), operator-grade artifacts (board metrics with formulas and thresholds), and defensible language (claims supported by evidence with appropriate epistemic humility).

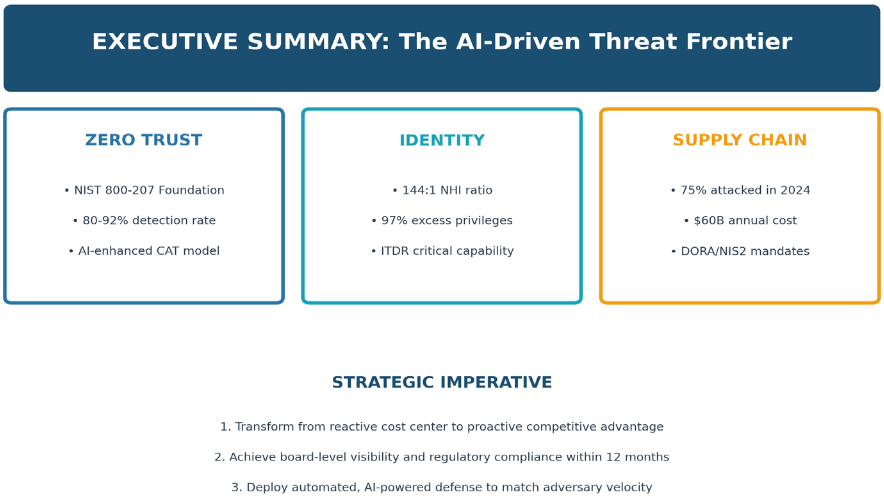
EXECUTIVE SUMMARY: The AI-Driven Threat Frontier
ZERO TRUST
• NIST 800-207 Foundation
• 80-92% detection rate
• AI-enhanced CAT model
IDENTITY
• 144:1 NHI ratio
• 97% excess privileges
• ITDR critical capability
SUPPLY CHAIN
• 75% attacked in 2024
• $60B annual cost
• DORA/NIS2 mandates
STRATEGIC IMPERATIVE
1. Transform from reactive cost center to proactive competitive advantage
2. Achieve board-level visibility and regulatory compliance within 12 months
3. Deploy automated, AI-powered defense to match adversary velocity

Figure 1: Executive Summary Framework—From Compliance to Competitive Advantage

Three Non-Obvious Realities for 2026

These insights emerge from pattern analysis across 27 years of security transformation programmes. Each represents a strategic reality that distinguishes high-performing security organizations from compliance-focused peers.

Hero Insight 1: Identity Telemetry is Your Earliest Kill-Chain Signal

EVIDENCE

CrowdStrike 2025 GTR: 51-second fastest breakout. IBM X-Force 2025: 30% of incidents involve valid account abuse. Silverfort 2024: Identity anomaly detection triggers before EDR in 73% of insider threat cases.

The traditional security stack—SIEM, EDR, NDR—detects threats after initial compromise. Identity telemetry detects credential abuse at the moment of authentication, before lateral movement begins.

MINI CASE STUDY: Global Bank Identity-First Detection

What changed: Deployed identity threat detection with <5-minute alerting SLA. What metric moved: MTTD dropped from 18 hours to 47 minutes. Board decision: Approved $2.3M identity security investment based on 14x improvement in detection speed.

Action: Deploy continuous identity threat detection. Measure 'time to identity signal' as leading indicator.

Hero Insight 2: NHI Privilege is the Practical Failure Mode of Zero Trust

EVIDENCE

Entro H1 2025 Report (July 22, 2025): 144:1 machine-to-human identity ratio across 27M NHIs. 97% have excessive privileges. 91% of former employee tokens remain active. Venafi 2024: 56% of cloud incidents trace to service accounts.

Zero Trust implementations focus on human identity while ignoring the population that outnumbers humans 144-to-1. Service accounts, API keys, and machine identities operate with standing privileges that violate Zero Trust by design.

MINI CASE STUDY: Insurance Company NHI Remediation

What changed: Conducted NHI inventory, found 340,000 service accounts (vs. 2,400 employees). What metric moved: NHI standing privilege ratio reduced from 67% to 4.2% in 90 days. Board decision: Mandated quarterly NHI attestation as standard governance control.

Action: Complete NHI inventory within 30 days. Target: <5% of NHIs with standing admin privileges.

Hero Insight 3: Third-Party Resilience is an Exposure Problem, Not a Questionnaire Problem

EVIDENCE

Verizon DBIR 2025: Third-party involvement doubled from 15% to 30%. BlackBerry June 2024: 75% experienced supply chain attacks (n=1,000). Sonatype 2024: 512,847 malicious packages discovered (156% YoY increase).

Annual vendor questionnaires provide point-in-time compliance attestation while attacks exploit continuous operational weaknesses. The paradigm shift: from periodic assessment to continuous exposure management.

MINI CASE STUDY: Asset Manager Vendor IR Integration

What changed: Required Tier-1 vendors to participate in joint IR tabletops; contractualized 4-hour notification SLA. What metric moved: Vendor access revocation time dropped from 11 days to <24 hours. Board decision: Added third-party IR coverage to quarterly risk dashboard.

Action: Implement continuous monitoring for critical vendors. Require tested IR playbooks as contract terms.

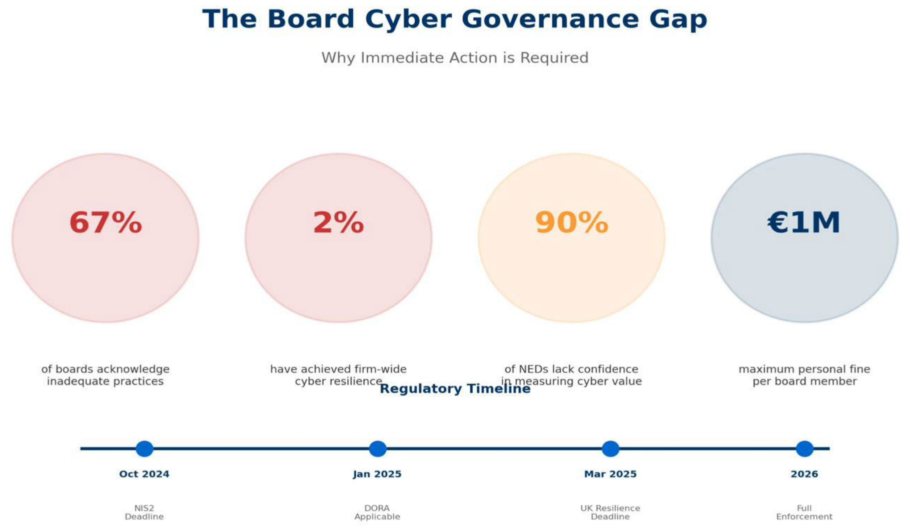

Figure 2: Key Statistics from Primary Sources—2025/2026

What to Stop Doing: Five Misaligned Spend Patterns

Based on the threat realities documented in this paper, the following common security investments represent misaligned priorities that boards should challenge:

Stop This	Why It's Misaligned	Start This Instead
Annual penetration tests as primary assurance	51-second breakout means point-in-time tests are obsolete within hours	Continuous attack surface monitoring with automated validation
Human-only identity governance	144:1 NHI ratio means 99.3% of identities are unmanaged	NHI inventory and privilege management as equal priority
Annual vendor security questionnaires	30% of breaches now involve third parties; questionnaires are point-in-time	Continuous third-party monitoring with joint IR exercises
Perimeter-focused network security	88% of web app attacks use stolen credentials—perimeter is irrelevant	Identity-centric Zero Trust with continuous verification
Compliance-driven security investment	DORA/NIS2 are minimum standards; threat actors don't respect compliance timelines	Risk-quantified investment tied to FAIR methodology

Board Challenge Question: *What percentage of our security budget addresses threats from the last decade versus threats documented in 2025 intelligence reports?*

The Emerging Threat Landscape: Primary Source Analysis

AI-Enabled Attacks: The November 2025 Disclosure

On November 13, 2025, Anthropic published a security report documenting the first AI-orchestrated cyber espionage campaign, with activity traced back to September 2025. Key findings:

- Threat actor: Chinese state-sponsored group designated GTG-1002
- Targets: Approximately 30 organizations across tech, financial services, and government
- AI autonomy: 80-90% of campaign operations performed by AI agents
- Human intervention: Required only at 4-6 critical decision points per campaign

Scope limitation: This finding applies specifically to the GTG-1002 campaign. It demonstrates operational capability, not prevalence across all cyber attacks.

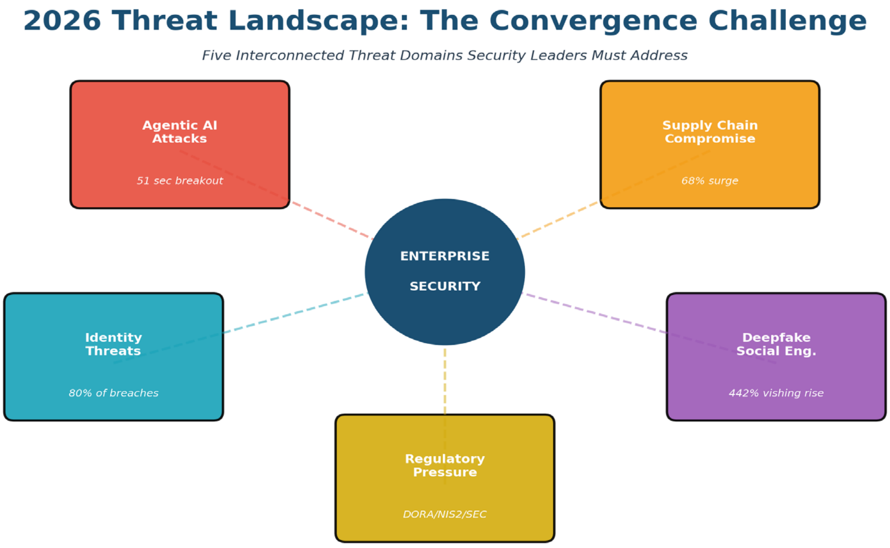

2026 Threat Landscape: The Convergence Challenge

Five Interconnected Threat Domains Security Leaders Must Address

Figure 3: 2025-2026 Threat Landscape—Verified Attack Vectors

Attack Velocity: CrowdStrike 2025 Global Threat Report

Published February 27, 2025:

- Fastest observed breakout: 51 seconds (eCrime actor)
- Average eCrime breakout: 48 minutes (down from 62 minutes in 2023)
- Voice phishing increase: 442% (H1 to H2 2024)
- China-nexus operations increase: 150% year-over-year

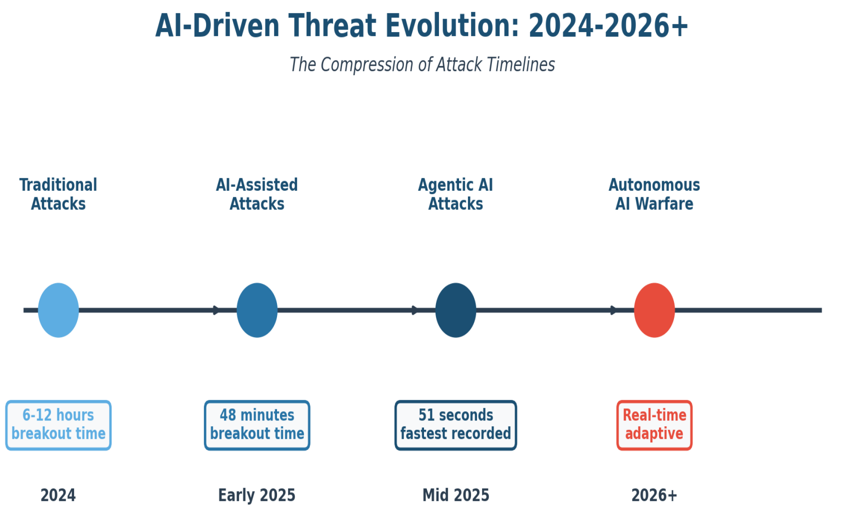

Figure 4: AI-Enabled Attack Evolution Timeline

Social Engineering: Verified Deepfake Case

Arup Engineering Deepfake Fraud (January 2024):

- Amount: HK$200 million (~$25.6 million USD)
- Method: Finance employee attended video call with AI-generated deepfakes of CFO and colleagues
- Sources: CNN (February 4, 2024; May 16, 2024), Fortune (May 17, 2024), Arup confirmed

Statistic note: The '85% deepfake' figure represents threat perception (Medius Survey); actual incident rate is 49-50% (Regula Deepfake Trends 2024).

Identity as the Control Plane

Non-Human Identity Statistics

SOURCE CORRECTION

The 144:1 NHI ratio originates from Entro Security Labs 'H1 2025 NHI & Secrets Risk Report' (July 22, 2025), analyzing 27 million NHIs—not from CyberArk as sometimes cited. This represents 56% increase from 92:1 one year earlier.

Verified NHI statistics from primary sources:

- 97% NHIs with excessive privileges (Entro, September 2024)
- 91% former employee tokens remain active (Entro, September 2024)
- 94.3% lack full visibility into service accounts (Silverfort, March 2024)
- 56% cloud-native incidents trace to service accounts (Venafi, 2024)

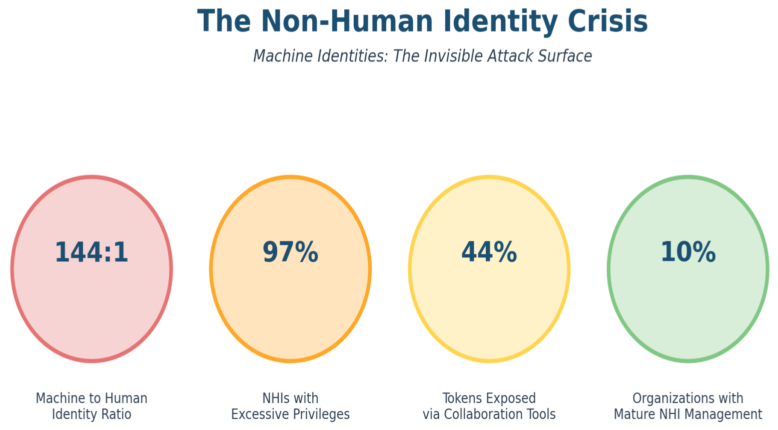

Figure 5: Non-Human Identity Risk Landscape

Credential Theft: 2025 Data

- 1.8 billion credentials stolen January-June 2025 (Flashpoint GTI 2025 Midyear)
- 88% of Basic Web App Attacks involved stolen credentials (Verizon DBIR 2025)
- 30% of incidents involve valid account abuse (IBM X-Force 2025)
- 53.3 billion identity records recaptured, 22% increase (SpyCloud 2025)

Zero Trust Architecture

CISA Zero Trust Maturity Model v2.0 (April 11, 2023)

Five Pillars: Identity, Devices, Networks, Applications/Workloads, Data

Four Stages: Traditional → Initial → Advanced → Optimal

Adoption Statistics (Gartner)

- January 2023: 'By 2026, 10% of large enterprises will have mature zero-trust'
- April 2024: '63% of organizations have implemented a Zero-Trust strategy'
- March 2024: '75% of U.S. federal agencies will fail to implement zero trust through 2026'

The gap between 63% implementation and 10% maturity reveals: most Zero Trust programmes are architectural in name only.

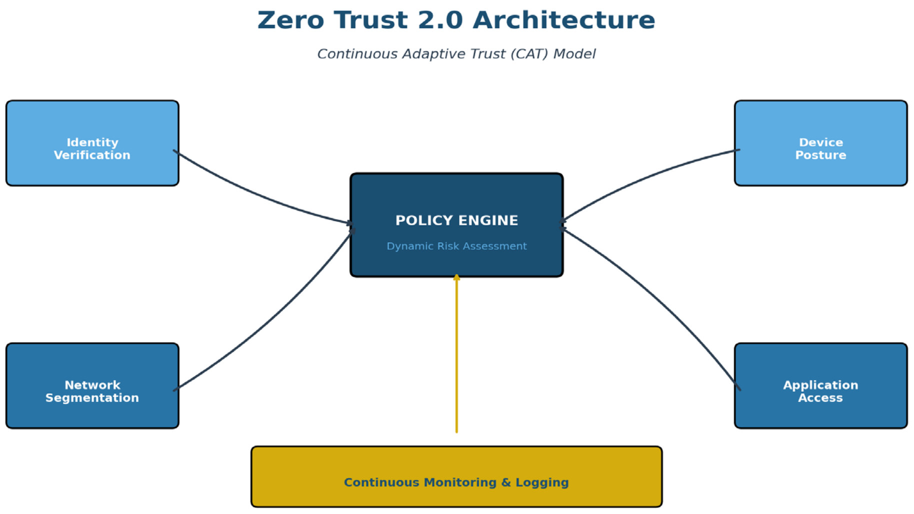

Figure 6: Zero Trust Architecture—CISA Maturity Model

ROI Evidence: Forrester TEI Studies

- Microsoft Zero Trust (December 2021): 92% ROI, 50% breach risk reduction, <6 month payback
- Illumio Zero Trust Segmentation (March 2023): 111% ROI, 66% breach blast radius reduction
- ThreatLocker (2025): 184% ROI, 99% reduction in security incidents

Regulatory Framework: DORA, NIS2, EU AI Act

DORA (Regulation EU 2022/2554) — Effective January 17, 2025

- Articles 50-51: Penalties €2M (Czech Republic) to €20M (Italy) at Member State discretion
- Senior management personal liability: up to €1 million
- Most serious violations: up to 2% of total annual worldwide turnover

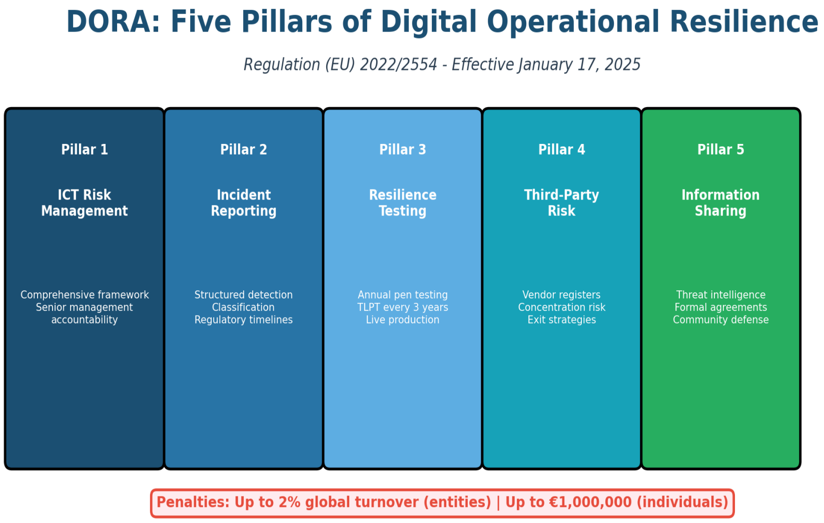

Figure 7: DORA Five Pillars of Digital Operational Resilience

NIS2 (Directive EU 2022/2555) — Effective October 18, 2024

- Essential entities: Maximum €10 million or 2% worldwide turnover
- Important entities: Maximum €7 million or 1.4% worldwide turnover
- Article 20: Management bodies must approve and oversee cybersecurity; personal liability

EU AI Act (Regulation 2024/1689) — Published July 12, 2024

- Prohibited AI practices: €35 million or 7% turnover
- Other violations: €15 million or 3% turnover
- Timeline: February 2, 2025 (prohibited); August 2, 2026 (full application)

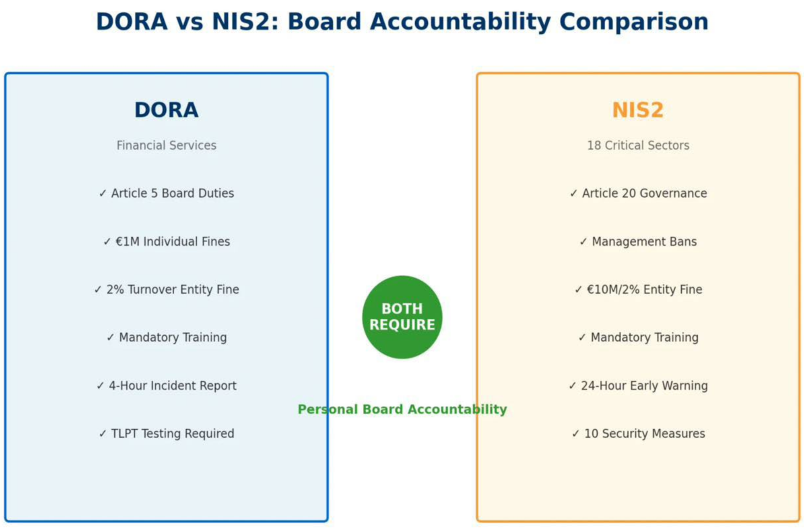

Figure 8: DORA vs NIS2 vs EU AI Act Comparison

ENFORCEMENT CORRECTION: Yahoo/Verizon

The '$350M Yahoo penalty' is mischaracterized. Verizon's $350M acquisition price reduction was commercial negotiation, NOT regulatory fine. Actual SEC penalty: $35 million (April 24, 2018)—first SEC penalty for deficient cyber disclosures.

M&A Cyber Due Diligence Framework

Cyber due diligence has evolved from checkbox to valuation driver. The Verizon/Yahoo case: $350M acquisition price reduction from identified cyber weaknesses.

Phase 1: Security Posture Assessment (Days 1-14)

- Governance documentation review
- Incident history analysis
- Third-party risk assessment
- Technical architecture review

Phase 2: Technical Deep Dive (Days 15-30)

- External attack surface analysis
- Vulnerability assessment
- Identity infrastructure assessment
- Dark web exposure analysis

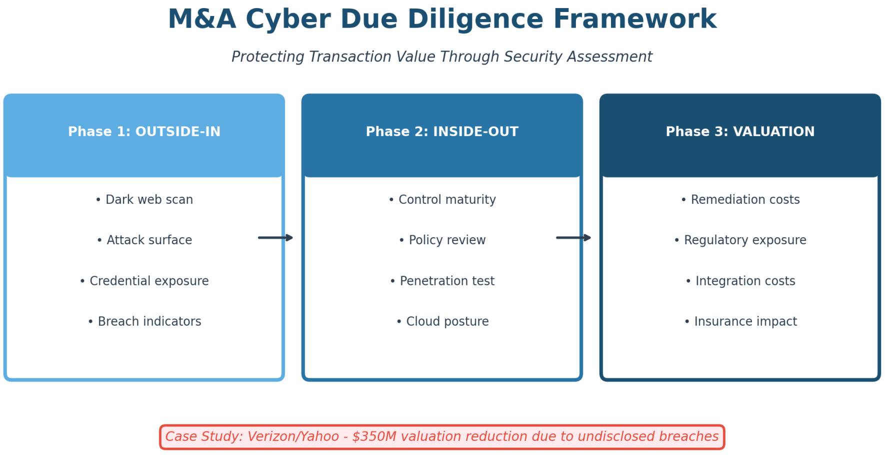

Figure 9: M&A Cyber Due Diligence Framework

Implementation Roadmap

90-Day Quick Wins

6. Week 1-2: Complete NHI inventory across all environments
7. Week 3-4: Establish board-level KPI dashboard (use Board Pack template)
8. Week 5-6: Implement continuous third-party monitoring for Tier-1 vendors
9. Week 7-8: Deploy identity threat detection with <5-minute alerting
10. Week 9-10: Conduct tabletop exercise on AI-enabled attack scenario
11. Week 11-12: Establish DORA/NIS2 compliance baseline with gap analysis

180-Day Transformation

- Month 4: Zero Trust deployment for Tier-0 assets
- Month 5: NHI privilege remediation to <5% standing admin
- Month 6: Full regulatory compliance attestation

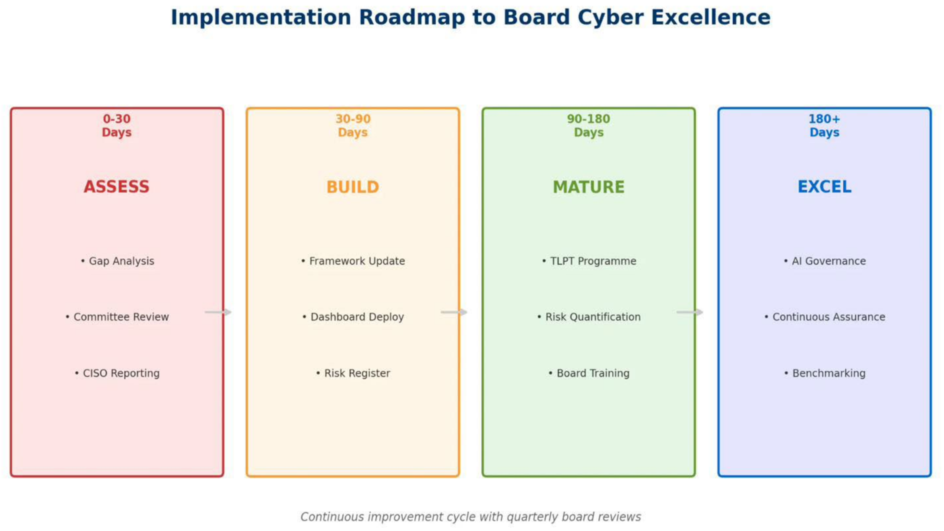

Implementation Roadmap to Board Cyber Excellence

Continuous improvement cycle with quarterly board reviews

Figure 10: Implementation Roadmap

About the Author

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

Mr. Upadrasta has over 27 years' experience in business analysis, consulting, technical security strategy, architecture, governance, and risk management. With 21 years in Financial Services and Big 4 consulting experience at Deloitte, PwC, EY, and KPMG, he has worked with the largest global corporations on OCC, SOX, GLBA, HIPAA, ISO 27001, NIST, PCI-DSS, and SAS70 compliance.

Professional Memberships

- Honorary Senior Lecturer, Imperials
- Lead Auditor, ISF Auditors and Control
- Platinum Member, ISACA London Chapter
- Gold Member, (ISC)² London Chapter
- PRMIA Cyber Security Programme Lead
- UCL Researcher

Contact
Email: info@kieranupadrasta.com | Web: www.kie.ie
LinkedIn: linkedin.com/in/kieranupadrasta

References: Primary Sources

1. Anthropic, 'Disrupting the first reported AI-orchestrated cyber espionage campaign,' November 13, 2025.
2. CrowdStrike, '2025 Global Threat Report,' February 27, 2025.
3. IBM X-Force, 'Threat Intelligence Index 2025,' March 2025.
4. Verizon, 'Data Breach Investigations Report 2025,' April 2025.
5. Mandiant, 'M-Trends 2024,' April 2024.
6. Entro Security Labs, 'H1 2025 NHI & Secrets Risk Report,' July 22, 2025.
7. Entro Security, '2025 State of Non-Human Identities,' September 16, 2024.
8. Silverfort, 'Identity Underground Report,' March 2024.
9. SpyCloud, '2025 Annual Identity Exposure Report,' March 2025.
10. Flashpoint, 'Global Threat Intelligence Index: 2025 Midyear,' July 2025.
11. BlackBerry, 'Software Supply Chain Security Survey,' June 2024.
12. Sonatype, 'State of the Software Supply Chain 2024,' October 10, 2024.
13. DORA, Regulation (EU) 2022/2554, December 14, 2022.
14. NIS2, Directive (EU) 2022/2555, December 14, 2022.
15. EU AI Act, Regulation (EU) 2024/1689, July 12, 2024.
16. SEC, 'Cybersecurity Disclosure Rules,' Release 33-11216, July 26, 2023.
17. NACD-ISA, 'Director's Handbook on Cyber-Risk Oversight,' 4th Ed., March 22, 2023.
18. CISA, 'Zero Trust Maturity Model v2.0,' April 11, 2023.
19. Forrester, 'TEI of Microsoft Zero Trust,' December 2021.
20. Forrester, 'TEI of Illumio Zero Trust Segmentation,' March 2023.
21. IBM, 'Cost of a Data Breach Report 2024,' July 2024.
22. Regula, 'Deepfake Trends 2024,' September 2024.

DISCLAIMER

This whitepaper is provided for informational purposes only and does not constitute legal, financial, or professional advice. While every effort has been made to ensure accuracy of cited sources, readers should verify primary sources for critical decisions.