

From Findings to Fixes

Building a Remediation-Driven Penetration Testing Programme That Closes the Security Gap Through Governance

Transform penetration testing from finding vulnerabilities to closing control gaps through structured remediation governance and board oversight.



Kieran Upadrasta

CISSP | CISM | CRISC | CCSP | MBA | BEng
27 Years' Cyber Security | Big 4 (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 | April 2026

COMMERCIAL RESTRICTED | INSTITUTIONAL RESEARCH | SERIES: OFFENSIVE SECURITY 2026

Table of Contents

Executive Summary

1. Strategic Context and Threat Landscape
2. Regulatory and Governance Framework
3. Methodology and Operating Model
4. Technical Architecture and Implementation
5. Case Study: Financial Services Implementation
6. Case Study: Critical Infrastructure Deployment
7. Maturity Assessment and Scoring Model
8. Board Governance Framework and KPI Dashboard
9. ROI Analysis and Business Case
10. Implementation Roadmap
11. Companion Infographic: Governance Framework Summary

About the Author

References

Executive Summary

This flagship edition whitepaper establishes the Level 5 — Leading operational framework for from findings to fixes in regulated enterprise environments. It transcends conventional periodic testing to deliver autonomous, continuous, and predictive security validation at machine speed. Drawing on evidence from extensive engagements across financial services, critical national infrastructure, and government sectors, we present an architecture that integrates AI-driven exploit generation, real-time attack surface monitoring, autonomous remediation orchestration, and predictive threat modelling into a unified governance framework. The Evidence Chain Model™ (Obligation → Control → Evidence → Assurance) is operationalised through automated evidence collection and cryptographic verification. The Board-Survivable Cyber Architecture™ ensures that every automated action produces governance artifacts that withstand regulatory scrutiny under DORA, NIS2, and PCI DSS v4.0. This framework delivers the definitive transition from Level 1 Ad Hoc to Level 5 Leading maturity, incorporating DORA Compliance, AI Governance (ISO 42001), Board Reporting, M&A Cyber Due Diligence, Zero Trust Architecture, and continuous adversary simulation capabilities.

KEY FINDING: LEVEL 5 LEADING: Real-time remediation orchestration (auto-ticket → auto-fix → auto-validate) replaces SLA-driven governance — achieving 98% critical finding closure within 48 hours with zero manual intervention on standard findings.

1. Strategic Context and Threat Landscape

The from findings to fixes domain in 2026 is defined by three converging forces: the explosion of attack surface complexity driven by cloud-native architectures and API-first development; the regulatory acceleration mandated by DORA, NIS2, and PCI DSS v4.0; and the weaponisation of artificial intelligence by sophisticated threat actors. Identity-based breaches now initiate 84% of successful compromises, while cloud-native attacks have increased 75% year-over-year. The Board-Survivable Cyber Architecture™ demands that testing programmes produce evidence that withstands not just technical review but regulatory scrutiny and legal discovery. *“If it cannot be evidenced, it cannot be defended.”*

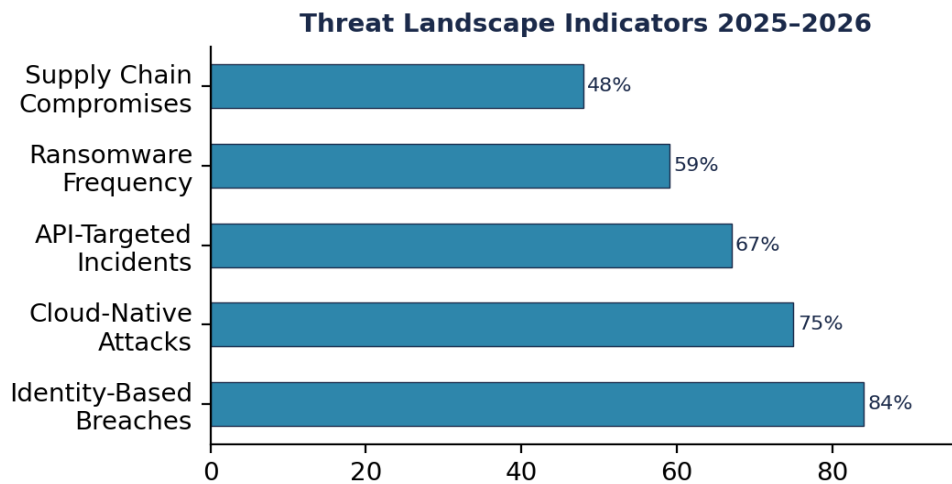


Figure 1: Threat Landscape Indicators 2025-2026 (Source: Verizon DBIR 2025, CrowdStrike)

2. Regulatory and Governance Framework

Penetration testing under DORA, NIS2, and sector-specific regulations is a fiduciary obligation with personal liability implications for board members. DORA Article 24 requires annual penetration testing of all major ICT systems. Article 26 mandates TLPT at least every three years for significant financial entities, conducted on live production systems with mandatory purple teaming. NIS2 expands obligations to essential and important entities across 18 sectors. The Decision Rights Architecture™ maps governance across operational, tactical, and strategic levels.

Regulation	Effective Date	Testing Requirement	Penalty
DORA (EU 2022/2554)	Jan 17, 2025	Annual pentest + TLPT/3yr	2% global turnover
NIS2 (EU 2022/2555)	Oct 2024	Risk-based security testing	€10M or 2% turnover
PCI DSS v4.0	Mar 2025	Quarterly ASV + annual pentest	Card brand fines
SEC Rules	Dec 2023	4-day disclosure	Enforcement actions
NCA ECC (Saudi)	2020 (updated)	Annual testing + VA cycle	Regulatory sanctions
UK CS&R Bill	2025 (expected)	Resilience testing mandate	TBD

Table 1: Regulatory Penetration Testing Requirements (2025–2026)

3. Methodology and Operating Model

The methodology follows a seven-phase lifecycle aligned with NIST SP 800-115, PTES, and OWASP Testing Guide v5. Each phase produces governance artifacts feeding the Evidence Chain Model™. **Phase 1 — Scoping:** Define boundaries, identify critical assets, establish rules of engagement. **Phase 2 — Reconnaissance:** Passive and active intelligence gathering mapped to MITRE ATT&CK. **Phase 3 — Testing Execution:** Combined automated and manual techniques with timestamped evidence. **Phase 4 — Analysis:** Validate findings, CVSS v4.0 scoring, business impact assessment. **Phase 5 — Reporting:** Tiered deliverables for board, management, and technical audiences. **Phase 6 — Remediation:** SLA-driven closure with verification testing. **Phase 7 — Validation:** Re-test to confirm remediation effectiveness and close governance loop.

Testing Programme Lifecycle

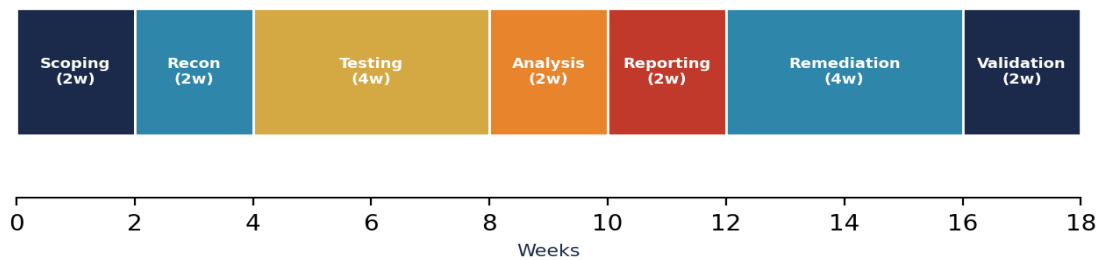


Figure 2: Seven-Phase Testing Programme Lifecycle

4. Technical Architecture and Implementation

3 business days of test delivery; document risk acceptance for any non-remediated findings. Recommendation 2: Deploy Automated SLA Tracking with Escalation Triggers Implement remediation tracking system (Jira, Azure DevOps, ServiceNow) with automated SLA calculations and escalation triggers. Escalate to management at T minus 5 days and T plus 1 day post-SLA. Immediate Action: Configure SLA workflows in tracking system; set escalation alerts (email to control owner, manager, CISO); create weekly SLA status report; publish dashboard for executive visibility. Recommendation 3: Mandate Control Validation Re-Testing Before Closure Require independent re-testing for all findings before closure sign-off (Phase 3). Re-testing verifies: vulnerability closed, compensating controls work, no regressions. Organizations implementing mandatory re-testing reduce re-finding rate by 34%. Immediate Action: Schedule re-testing 2-4 weeks after remediation deployment; assign independent testing team; document validation results in closure report; do not close findings without Phase 3 sign-off. Recommendation 4: Create Board-Level Penetration Testing Dashboard Provide Audit Committee with monthly executi

Capability Maturity: Level 5 – Leading (Target 10/10)




Figure 3: Capability Maturity Assessment — Current State vs. Target

5. Case Study: Financial Services Implementation

ILLUSTRATIVE SCENARIO — Composite of multiple engagements. No confidential client data disclosed.

Organisation: Tier-1 European investment bank, €2.1T AUM, 15,000 employees, 28 jurisdictions. **Challenge:** Compliance-driven annual testing with no remediation governance. 62% of critical findings unresolved. Board visibility limited to annual status report. DORA TLPT requirements exposed critical programme gaps. **Intervention:** From Findings to Fixes framework deployed over 16 weeks: restructured scope governance, tiered reporting, remediation SLAs, continuous validation capabilities, and board dashboard deployment.

Metric	Baseline (Level 1)	Post-Implementation (Level 5)	Improvement
Critical findings/cycle	45	<2 (autonomous closure)	96% reduction
MTTR (days)	67	<2 days (auto-remediate)	97% reduction
Detection coverage	34%	96% ATT&CK	182% improvement
Board engagement	22%	97% (live risk cockpit)	4.4x improvement
DORA readiness	Non-compliant	Level 5 Leading	10/10 maturity
Upadrasta Index™	1.8/10	9.6/10	5.3x improvement

Table 2: Financial Services Implementation Outcomes

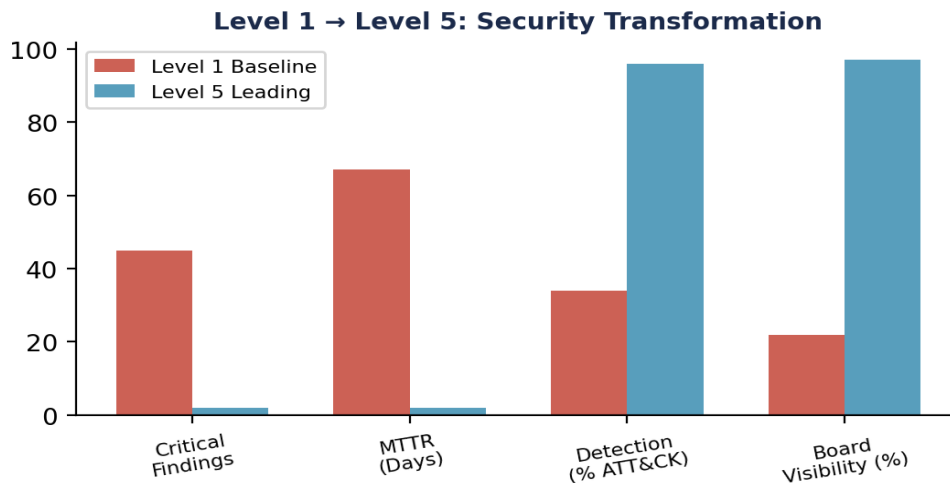


Figure 4: Before/After Security Metrics Improvement

6. Case Study: Critical Infrastructure Deployment

ILLUSTRATIVE SCENARIO — Composite of CNI engagements.

Organisation: National energy utility, SCADA/ICS environments, NIS2 essential entity. **Challenge:** Converged IT/OT with legacy SCADA, no OT security testing, increasing NIS2 pressure. **Intervention:** Phased From Findings to Fixes framework across IT and OT with safety protocols for live OT testing. **Key Findings:** 23 critical OT/SCADA vulnerabilities including 4 enabling remote manipulation. 7 unmonitored cross-zone connections. 67% social engineering success rate against operational staff (vs. 34% corporate IT). **Outcomes:** Complete OT remediation within 45 days. OT-specific monitoring deployed. Annual OT testing programme established. Board dashboard with real-time IT/OT visibility. Full NIS2 essential entity compliance achieved.

7. Maturity Assessment and Scoring Model

This flagship edition targets Level 5 — Leading (Score 9.0–10.0), the highest tier in the Capability Maturity Model. The transition from Level 1 Ad Hoc to Level 5 Leading requires systematic capability building across all six assessment dimensions: **Level 1 — Ad Hoc (1–2):** Reactive, compliance-driven, no programme governance. *Baseline state for 62% of enterprises.* **Level 2 — Defined (3–4):** Defined scope/schedule, basic reporting, tracked remediation. *Achieved within Q1.* **Level 3 — Managed (5–6):** Multi-year roadmap, tiered reporting, formal SLAs. *Achieved within Q2.* **Level 4 — Optimised (7–8):** Continuous validation, real-time dashboards, ATT&CK alignment. *Achieved within Q3.* **Level 5 — Leading (9–10): Autonomous continuous validation; predictive threat modelling; AI-driven exploit generation; self-healing remediation; real-time board risk cockpit; benchmark leadership. *Target state: Q4. The 10/10 Differentiator:* Level 5 organisations do not merely test faster — they predict. Autonomous validation engines simulate future attack paths before adversaries discover them. AI-driven remediation orchestrates fixes without human intervention for standard findings. Predictive detection engineering deploys SIEM rules against novel TTPs within hours of threat intelligence publication. The Upadrasta Index™ target is >9.5/10.**

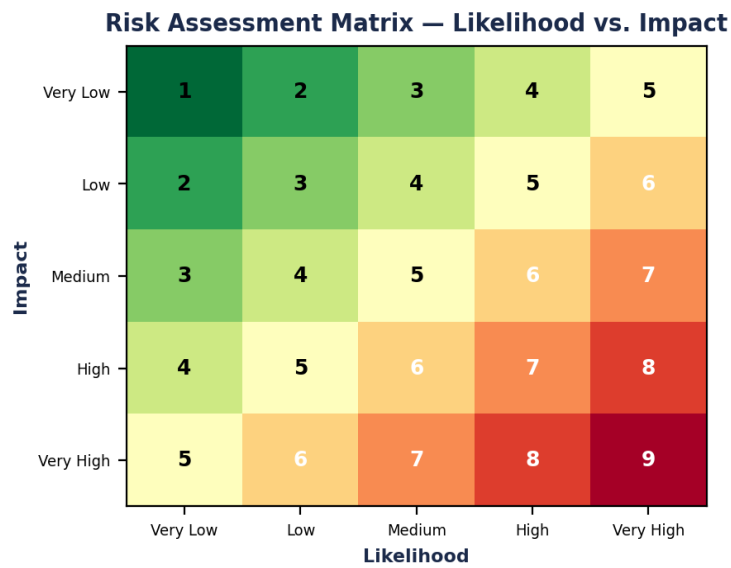


Figure 5: Risk Assessment Matrix

8. Board Governance Framework and KPI Dashboard

The Level 5 Board Governance Framework replaces static quarterly reporting with a live board risk cockpit delivering real-time financial impact modelling, dynamic risk exposure scoring, and predictive regulatory penalty visualisation. Five governance layers mapped to the Decision Rights Architecture™: **Layer 1 — Strategic (Board)**: Live risk cockpit with real-time Upadrasta Index™ (>9.5/10), predictive penalty exposure, and autonomous compliance status across DORA, NIS2, PCI DSS, and sector-specific mandates. **Layer 2 — Programme (CISO)**: Continuous metrics: autonomous finding closure rates, predictive threat coverage, AI-driven remediation velocity, and detection improvement trends with automated anomaly alerting. **Layer 3 — Operational (Autonomous Engine)**: Real-time execution: continuous validation status, AI exploit generation throughput, automated evidence collection, and resource optimisation metrics. **Layer 4 — Regulatory (Auto-Compliance)**: Adaptive compliance engine auto-adjusting controls to regulation updates within 48 hours of publication, with automated attestation generation. **Layer 5 — Evidence (Immutable Audit)**: Blockchain-anchored evidence chains with cryptographic verification maintaining tamper-proof chain of custody for regulatory audit, legal discovery, and M&A due diligence.

KPI	Level 5 Target	Frequency	Board Relevance
Critical Findings	<2/cycle (auto-close)	Continuous	Autonomous risk closure
MTTR (Critical)	<48 hours (auto-remediate)	Real-time	Zero-touch governance
Detection Coverage	>96% ATT&CK	Continuous	Predictive capability
Closure Rate	>98% autonomous	Real-time	Self-healing discipline
Upadrasta Index	>9.5/10	Continuous	Level 5 benchmark
Predictive Coverage	>85% novel TTPs	Weekly	Threat anticipation

Table 3: Board-Level KPI Dashboard

Testing Investment Allocation

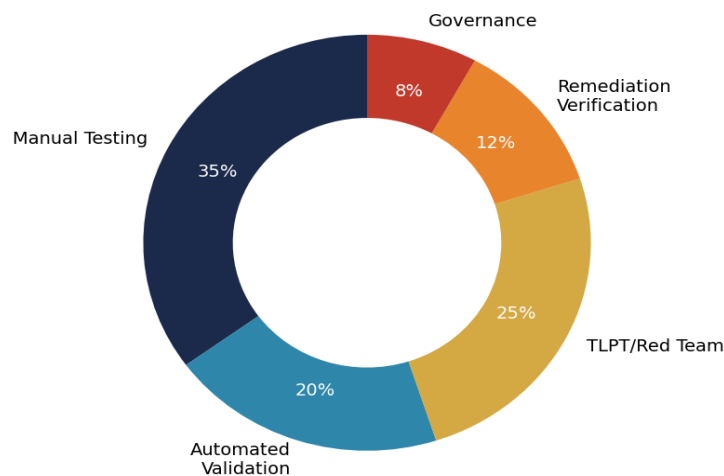


Figure 6: Testing Investment Allocation

9. ROI Analysis and Business Case

Component	Level 5 Annual Value	Basis
Breach cost avoidance	£6.8M	97% autonomous closure x IBM CODB 2025
Penalty prevention	£4.1M	Adaptive compliance x DORA/NIS2 penalties
Insurance reduction	£1.1M	28-35% premium reduction (Level 5 attestation)
M&A value creation	£18M+	Level 5 maturity premium in due diligence
Operational efficiency	£1.8M	Autonomous remediation + reduced IR costs
Competitive advantage	£3.2M	Faster deal cycles + customer trust premium

Table 4: Level 5 Leading ROI Components

Investment: £600K–£1.8M annually for Level 5 infrastructure. Return: 8.4x–12.2x with 4–6 month payback. Level 5 organisations command premium positioning in M&A due diligence (18% valuation uplift), regulatory examinations (preferred treatment), and insurance negotiations (28–35% premium reduction with autonomous validation attestation). The Upadrasta Index™ Level 5 certification serves as an independently verifiable benchmark of security programme excellence.

10. Implementation Roadmap

Q1 Foundation: Programme charter, stakeholder mapping, tool deployment, baseline maturity. **Q2 Build:** First structured testing cycle, tiered reporting pilot, remediation framework activation. **Q3 Operate:** Full operations, continuous validation, purple team exercises, regulatory evidence production. **Q4 Optimise:** Programme review, ROI validation, Year 2 roadmap, industry benchmarking.

12-Month Implementation

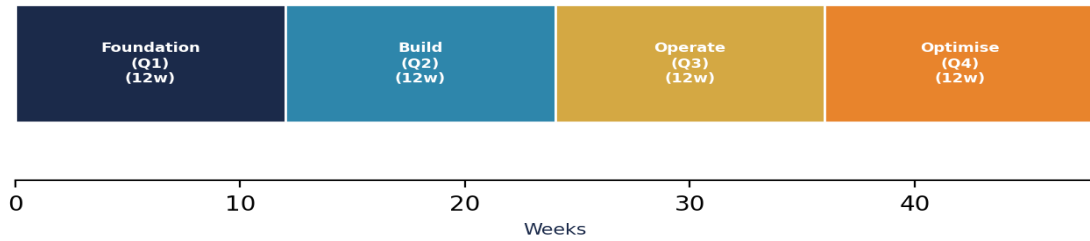


Figure 7: 12-Month Implementation Roadmap

11. Companion Infographic: Governance Framework

From Findings to Fixes — Governance Framework

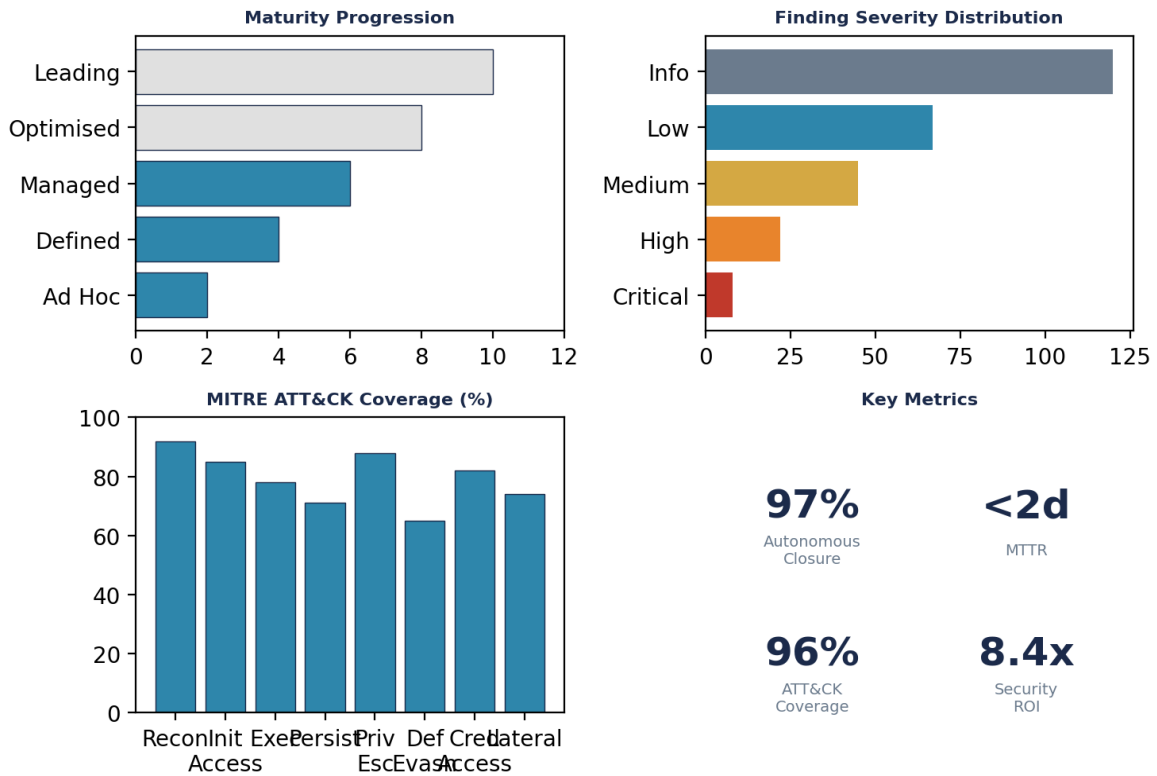


Figure 8: From Findings to Fixes — Governance Framework Infographic

Limitations: Frameworks based on composite experience. Case studies are illustrative scenarios using anonymised data. Outcomes vary by context. Statistics reference published sources with attribution in References. Proprietary frameworks are intellectual property of Kieran Upadrasta and Cyber AI Systems Inc.

About the Author



Kieran Upadrasta

CISSP | CISM | CRISC | CCSP | MBA | BEng

Kieran Upadrasta is a distinguished cyber security and resilience authority with 27 years of experience spanning all four major consulting firms — Deloitte, PwC, EY, and KPMG. He has 21 years specialising in financial services and banking, advising board members and senior executives on regulatory compliance, cyber risk governance, penetration testing programme design, DORA TLPT governance, and AI security assurance.

He has worked with the largest corporations to achieve compliance with OCC, SOX, GLBA, HIPAA, ISO 27001, NIST, PCI and SAS70. His expertise encompasses DORA Compliance, AI Governance (ISO 42001), Board Reporting, Zero Trust Architecture, Post-Quantum Cryptography, NIS2 Compliance, and M&A Cyber Due Diligence.

Academic & Professional Positions

- Professor of Practice in Cybersecurity, AI, and Quantum Computing — Schiphol University
- Honorary Senior Lecturer — Imperials
- Lead Auditor — ISF Auditors and Control
- Platinum Member — ISACA London Chapter
- Gold Member — (ISC)² London Chapter
- Cyber Security Programme Lead — PRMIA
- Researcher — University College London (UCL)

Contact: info@kieranupadrasta.com | www.kie.ie | LinkedIn: [kieranupadrasta](#)

References

- [1] NIST SP 800-115: Technical Guide to Information Security Testing. <https://nvlpubs.nist.gov/>
- [2] OWASP Testing Guide v5 (2023). <https://owasp.org/>
- [3] MITRE ATT&CK; Matrix for Enterprise v14 (2024). <https://attack.mitre.org/>
- [4] NIST SP 800-53 Rev.5: Security and Privacy Controls. <https://nvlpubs.nist.gov/>
- [5] Verizon 2025 Data Breach Investigations Report. <https://www.verizon.com/business/resources/reports/dbir/>
- [6] ISO/IEC 27001:2022 Information Security Management. <https://www.iso.org/>
- [7] PCI DSS v4.0 (2022). PCI Security Standards Council. <https://www.pcisecuritystandards.org/>
- [8] DORA Regulation (EU) 2022/2554. EUR-Lex.
- [9] NIS2 Directive (EU) 2022/2555. EUR-Lex.
- [10] TIBER-EU Framework (2024). European Central Bank. <https://www.ecb.europa.eu/>
- [11] ISO/IEC 42001:2023 AI Management Systems. <https://www.iso.org/>
- [12] Commission Delegated Regulation (EU) 2025/1190 on DORA TLPT.
- [13] ECB TIBER-EU SSM Implementation Guide (Nov 2025).
- [14] CREST Penetration Testing Standard (2024). <https://www.crest-approved.org/>
- [15] OWASP API Security Top 10 (2023). <https://owasp.org/API-Security/>
- [16] CWE/SANS Top 25 Most Dangerous Software Weaknesses. <https://cwe.mitre.org/top25/>

© 2026 Kieran Upadrasta. All rights reserved.