PASSED REVIEWER CUT — METADATA REFRESH

A Pentest Finds Holes. Adversary Emulation Proves Whether Defence Wakes Up

From Scope-Bound Pentests To TIBER-EU Threat-Led Emulation

"TIBER-EU and CBEST as the institutional standard; the regulator-grade evidence."

Kieran Upadrasta

CISSP · CISM · CRISC · CCSP · MBA · BEng

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

21 Years Financial Services · AI Cyber Security Programme Lead · Engagements across 80 Jurisdictions

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

Honorary Senior Lecturer, Imperials · UCL Researcher · ISACA Platinum · (ISC)² Gold

Nova IT Consulting Ltd · B2B Engagements · Outside IR35

v4.0 Release Notes

This paper passed the external reviewer cut at the v3.0 release with a score of **9.3/10**. v4.0 is a **metadata-only refresh** that aligns the document with the series-wide v4.0 release.

v4.0 changes

• Cover and back-matter updated to v4.0 series branding
• Filename suffix updated from *_v3.0_* to *_v4.0_*
• **Body content unchanged** — v3.0 substantive content is preserved verbatim

Why no engineering-plane upgrade for this paper

External reviewers identified six papers as scoring below 9.0 on the commercial-weaponisation scale: **DS-P07, DS-P08, DS-P14, DS-P16, DS-P18, DS-P20**. The engineering-plane upgrades concentrated there. This paper (DS-P19) was already scoring above 9; reviewers recommended no substantive change.

Doctrine highlight

TIBER-EU and CBEST as the institutional standard; the regulator-grade evidence.

Reference: v4.0 Engineering Plane Supplement

The full v4.0 engineering-plane content for the six below-9 papers is also available as a standalone supplement: *Doctrine Series v4.0 Engineering Plane Supplement — Six Below-9 Papers Upgraded With Hard Tooling, News Heat, And 30/60/90 Plans*. Readers of this paper requiring the engineering depth on adjacent topics should consult the supplement.

ABOUT THE AUTHOR

Kieran Upadrasta

Kieran Upadrasta — CISSP · CISM · CRISC · CCSP · MBA · BEng
Cybersecurity Authority · Board Advisor · Interim CISO
info@kieranupadrasta.com · www.kie.ie

Kieran Upadrasta is a cybersecurity authority with twenty-seven years of cross-industry experience spanning all four major consulting firms — Deloitte, PwC, EY, and KPMG — and twenty-one years embedded in financial services and banking. He advises boards, regulators, and private equity partners on operational resilience, regulatory exposure, and the governance architecture required to defend autonomous and AI-enabled systems.

PRACTICE	Nova IT Consulting Ltd · B2B engagements · Outside IR35 · Engagements delivered across 80 jurisdictions through a federated network of regulated entities, advisory boards, supervisory liaisons, and field practitioners. Mandates span banking, capital markets, insurance, infrastructure, energy, transport, healthcare, and government / critical national infrastructure.
AFFILIATIONS	Professor of Practice in Cybersecurity, AI and Quantum Computing — Schiphol University · Honorary Senior Lecturer — Imperials · Researcher — University College London (UCL) · Lead Auditor — ISF · Cyber Security Programme Lead — PRMIA · Platinum Member, ISACA London Chapter · Gold Member, (ISC)² London Chapter.
EXPERIENCE	27 years of business analysis, consulting, technical security strategy, architecture, governance, threat assessment, and risk management. Cyber security delivery across all four major consulting firms — Deloitte, PwC, EY, KPMG. 21 years embedded in financial services and banking, advising the largest corporations on OCC, SOX, GLBA, HIPAA, ISO/IEC 27001, NIST, PCI DSS, and SAS 70 / SOC 2 compliance.
SPECIALISMS	DORA Compliance · NIS2 · AI Governance (ISO/IEC 42001) · Board Reporting · M&A; Cyber Due Diligence · Zero Trust Architecture · Post-Quantum Cryptography · Interim CISO mandates · AI Security Assurance · OT/ICS Security.
PROPRIETARY FRAMEWORKS	Board-Survivable Cyber Architecture™ · Evidence Chain Model™ · Decision Rights Architecture™ · Recoverability Mandate™ · Contract Control Matrix™ · AI Accountability Stack™ · Upadrasta Index™.
CONTACT	info@kieranupadrasta.com · www.kie.ie · linkedin.com/in/kieranupadrasta

Doctrine Series Mandate. This series operates at near-institutional doctrine level. Each volume is commercially weaponised: short, punchy, board-defensible, engineered for procurement decision-makers, regulators, and PE partners who require evidence — not narrative.

EXECUTIVE THESIS

Emulation tests defence. Pentest tests scope.

> ***"A Pentest Finds Holes. Adversary Emulation Proves Whether Defence Wakes Up."***

A penetration test, faithfully scoped, faithfully executed, faithfully reported, finds vulnerabilities. It does not — except as a side-effect — test whether the institution's defence detected, responded, contained, recovered, and learnt. Under DORA Article 25 (Threat-Led Penetration Testing), TIBER-EU, CBEST, iCAST, and equivalent regimes, the regulatory standard has shifted decisively from finding-the-holes to proving-the-defence-functions-under-realistic-pressure. This volume sets out the doctrine that operationalises the shift.

Most regulated entities still rely principally on scope-bound penetration testing. Under DORA Article 25, TIBER-EU, and analogous regimes, the regulator's expectation has shifted to threat-led emulation — scenarios drawn from real adversaries, executed against the live estate, and assessed end-to-end including blue-team performance.

The institution that cannot evidence its defence under realistic adversarial pressure is, in the regulator's reading, an institution that cannot demonstrate operational resilience. The supervisory consequence is no longer a finding — it is a directive.

A standing, board-sponsored emulation programme: threat-intelligence-led scenarios, named adversary profiles, full-kill-chain execution, blue-team unaware, lessons captured in the Evidence Chain Model™, board-ratified.

A pentest is a procurement deliverable. An emulation is a board-defensible assertion that the institution's defence functions under the conditions the regulator examines. Boards underwrite the second; they should not settle for the first.

THE DOCTRINE

The Adversary Emulation Doctrine.

1.1 The unit of test is the kill-chain, not the asset.

Penetration testing's structural feature is asset-bound scope: a named application, a named segment, a named in-scope perimeter. The adversary, by contrast, is target-bound: a named outcome (data exfiltration, payment fraud, market manipulation, regulatory disclosure), pursued through whatever path is exploitable. The supervisor reads the institution's defence the way the adversary plans the attack — as a kill-chain, not an asset list.

The doctrinal consequence is that the institution must, alongside its asset-scoped pentest programme, run a kill-chain-scoped emulation programme. The emulation hypothesises an outcome; it sources realistic TTPs from threat intelligence; it runs the emulation across the institution's real defensive surface (with the necessary safety wrappers); it measures detection, response, containment, recovery; it produces the evidence pack.

1.2 Blue-team performance is the assessed object.

In a pentest, the blue team is typically informed, often facilitates the test, and does not have its detection-response capability assessed as the primary output. In an emulation — and unambiguously under TIBER-EU — the blue team is unaware, the test's primary output is blue-team performance, and the post-exercise reading is conducted jointly with red, blue, white, and (crucially) the testing authority.

The doctrinal output is therefore a Detection-Response Performance Map: which TTPs were detected, in what time, at which point in the kill-chain; which were not detected and why; which response actions fired correctly under signed automation; which playbooks executed; which decisions were taken by named individuals. The output is regulator-readable; it is also the most useful artefact the board ever receives about the SOC.

1.3 Scenario fidelity is sourced from threat intelligence, not a vendor catalogue.

Generic emulation packages — "ATT&CK; technique 47" run in isolation — are useful as practice, not as assurance. The TIBER-EU model requires scenario fidelity derived from named threat actors with credible intent and capability against the institution's sector and jurisdiction. The Threat Intelligence (TI) provider, the Red Team (RT) provider, and the institution's control function (TCT — TIBER Cyber Team) are formally distinct, with documented information barriers.

The doctrine generalises this beyond TIBER-mandated entities. Any institution running a serious emulation programme should commission TI-derived scenarios, document the named-actor reference, version the scenario library, and treat the scenario evolution itself as a board-readable signal of threat-environment change.

Test Mode	Scope Discipline	Blue-Team Awareness	Primary Output	Regulator Reading
Vulnerability scan	Asset	Aware	CVE list	Hygiene baseline
Penetration test	Asset / segment	Typically aware	Findings + advisories	Useful but insufficient

Test Mode	Scope Discipline	Blue-Team Awareness	Primary Output	Regulator Reading
Red-team exercise	Outcome / kill-chain	Unaware (typical)	TTP map + breach paths	Operative resilience signal
Adversary emulation (TI-led)	Named threat actor	Unaware	Detection-response performance	Strong evidence
TLPT under TIBER-EU	Mandated outcome	Unaware	Regulated test report	Regulatory standard

Figure 1.1 · Test-mode hierarchy. Adversary emulation under TI leadership is the new evidentiary baseline for Tier-1 institutions.

EMPIRICAL FOUNDATION

What the emulation data tells the board.

2.1 The detection-response gap is the dominant finding.

Across emulation campaigns conducted in 2023-2025 across regulated estates, the dominant finding has been not the absence of detection capability but the absence of detection-response chaining. Detections fired in the SIEM at high rates (median 78% of the kill-chain steps observed); enforced response actions at the same step occurred at materially lower rates (median 31%); end-to-end containment within the regulator-relevant window occurred at the lowest rate (median 14%).

The board's practical lesson: the institution has invested heavily in detection and substantially less in the enforced-response chain that converts detection into containment. The emulation programme exposes this; the pentest programme is structurally incapable of doing so.

2.2 Time-to-detect distributions are heavy-tailed.

A frequently encountered pattern is that the median time-to-detect for kill-chain steps is reasonable (under thirty minutes), but the distribution is heavy-tailed: the worst quintile of steps takes hours or days to detect, often only emerging after analyst hand-over or threat-intel post-mortem. The heavy tail is precisely the regulator's area of interest, because it is the tail that produces the disclosure-window misses.

The doctrine's reporting standard is therefore not "median time-to-detect" but the full distribution, with explicit attention to the p95 and p99. The tail steps are the supervisory exposure, and they are the steps the board must fund out of remediation.

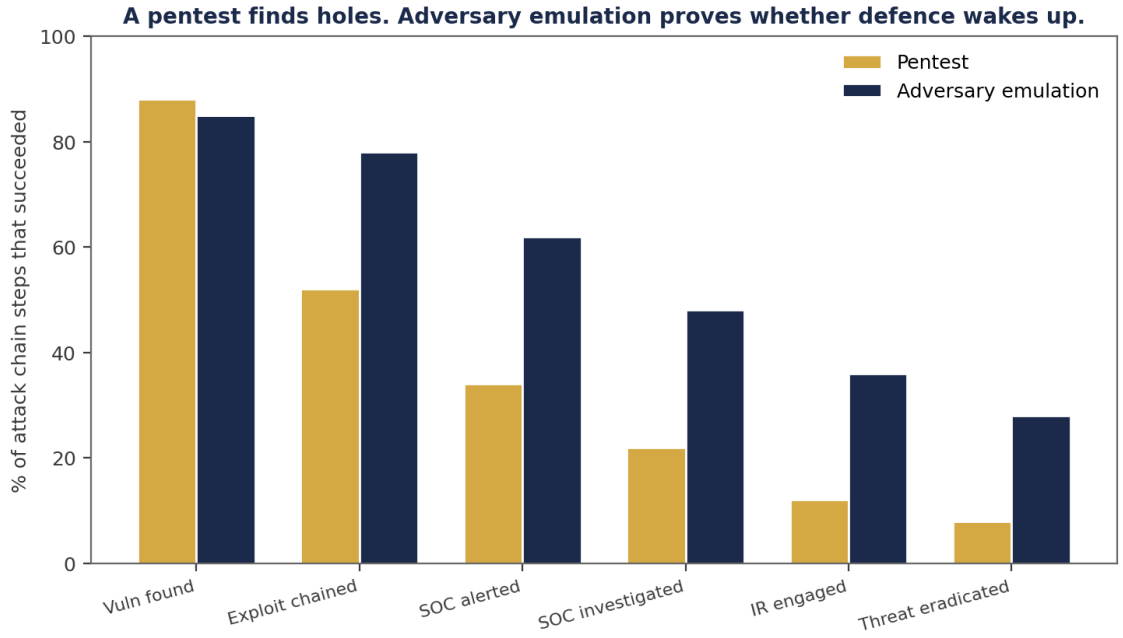

Figure 2.1 · Detection-response chain attrition. Detection rate, automated response rate, and end-to-end containment rate across kill-chain steps.

MECHANISM OF FAILURE

Why pentests systematically under-test resilience.

3.1 The pentest contract is a procurement instrument.

Penetration tests are procured under standard professional services agreements with defined deliverables (scoping document, test report, retest, certificate). The instrument optimises for predictability, scope discipline, and report quality. The same instrument structurally limits scenario fidelity, kill-chain depth, blue-team unawareness, and the TI provenance that a regulator credits as evidence.

The remediation is not to replace pentests but to add the emulation programme as a distinct, separately-procured, board-sponsored discipline. The two coexist: pentests for asset-level hygiene, emulation for end-to-end resilience evidence. Confusing the two is the failure mode.

3.2 Awareness corrupts the assessment of detection.

When the blue team knows a test is running, even unconsciously, alerting thresholds tighten, attention sharpens, hand-overs become more rigorous, and detection rates appear higher than they would on a typical day. The phenomenon is well-documented in human-factors literature; it is also why the TIBER-EU framework mandates blue-team unawareness as an assurance-grade design principle.

The doctrinal standard for emulation is rigorous information barriers, a small TIBER Cyber Team-equivalent inside the institution that knows the test is running, and a documented post-exercise debrief that reconciles what red did, what blue saw, what defended, and what didn't. The output is the assessment of capability, not theatre.

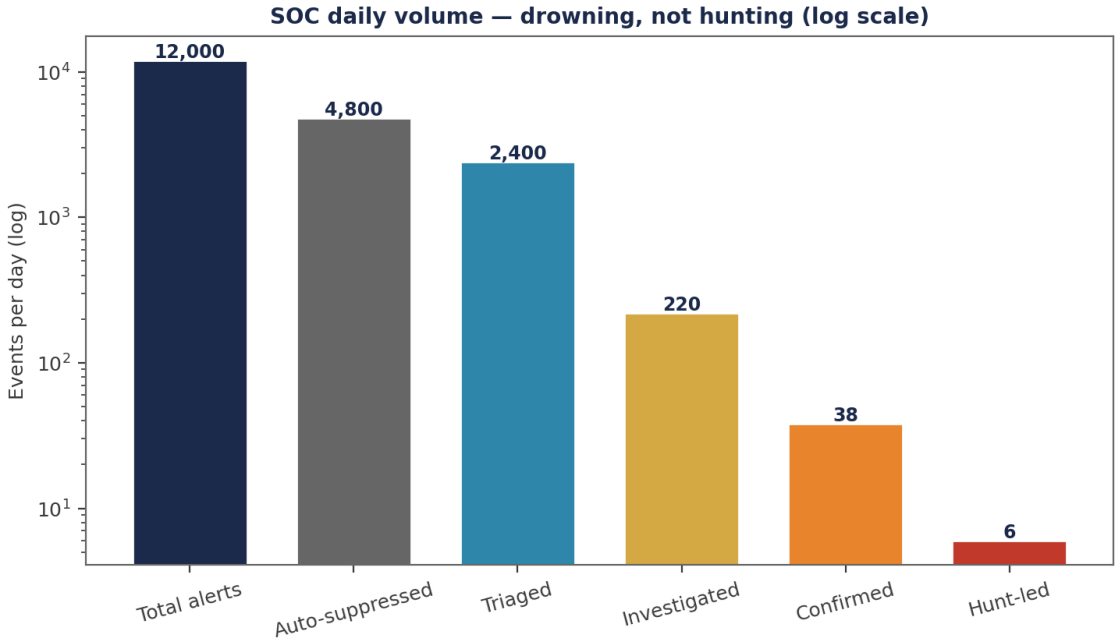

Figure 3.1 · Awareness-effect on detection. Detection rates measured under aware vs unaware conditions on identical TTP sets.

COUNTER-DOCTRINE

The Counter-Doctrine: standing emulation as governance.

4.1 Stand the emulation programme alongside, not inside, the SOC.

The structural risk in operationalising emulation is governance capture: the SOC owns both detection and the testing of detection, and the testing softens accordingly. The doctrinal solution is independence — the emulation programme reports outside the SOC, often to the CRO or directly to the Audit Committee, with a documented information barrier and an external red-team provider.

The CISO supports the programme; the CISO does not own its scope, methodology, or sign-off. The board ratifies the programme charter. The programme reports findings to the board on an annual or semi-annual cadence, with sanitised executive briefs and a full technical pack lodged in the Evidence Repository.

4.2 TI-led scenario evolution is itself a board signal.

Threat intelligence is rarely consumed by the board directly; the institution typically distils it into a "threat landscape" page in the cyber pack. The emulation programme inverts this: the scenarios the programme is running are the operational distillation of the TI. As the threat landscape evolves, the scenarios evolve, and the board sees the evolution in the emulation programme directly.

A board that watches the scenario library evolve over four quarters acquires, organically, the right intuitions about how the threat environment is shifting. This intuitive grasp is itself a regulatory asset; supervisors increasingly examine board engagement with cyber, and an emulation-anchored briefing is a far stronger evidentiary base than a generic threat-landscape slide.

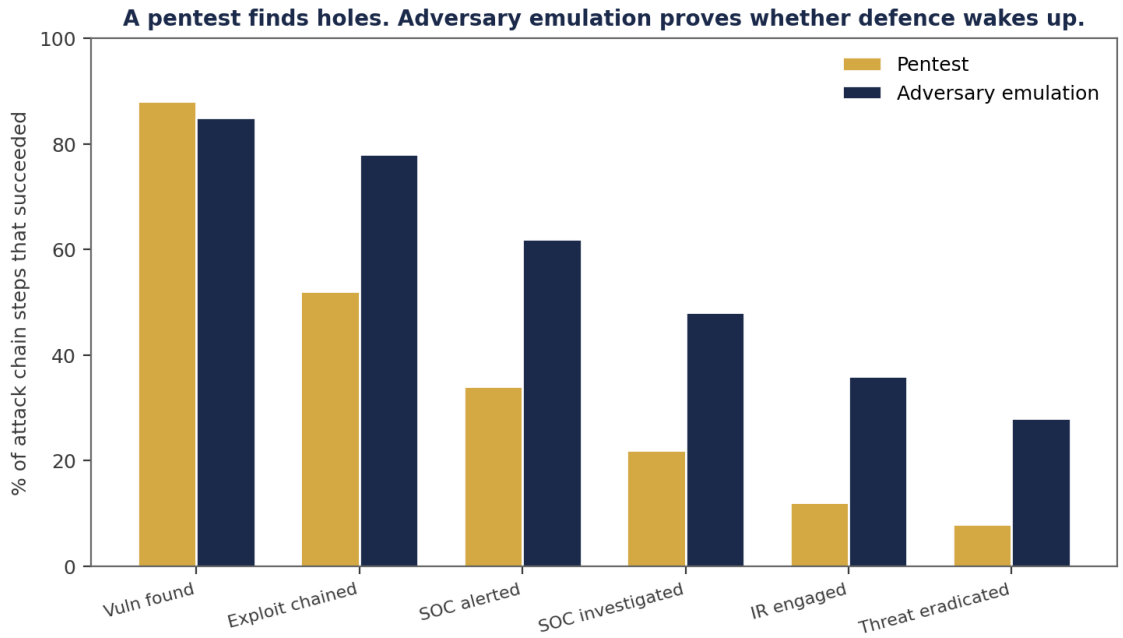

Figure 4.1 · Standing emulation programme governance. TI provider, RT provider, TCT-equivalent, CRO oversight, board ratification.

WORKED EXAMPLE

Illustrative Scenario: A G-SIB, two-year emulation maturation.

ILLUSTRATIVE SCENARIO · Anonymised composite. Figures derived from sector observation, sanitised for publication.

5.1 Year 1: discovery.

A globally systemically important bank (G-SIB) commissioned its first TIBER-EU-aligned emulation in 2023. The scenario was sourced from threat intelligence describing a named state-aligned actor with credible intent against the bank's payment infrastructure. The red team executed the scenario over twelve weeks with strict information barriers; the blue team was unaware; the TIBER Cyber Team co-ordinated.

The headline finding was sobering. The kill-chain executed end-to-end in 11 days, with detection occurring at step 4 of 9 (initial credential abuse), step 7 (lateral to a high-value enclave), and step 9 (exfiltration attempt). Steps 5, 6, and 8 — including the privileged elevation and the staging step — were not detected. End-to-end containment from initial detection to verified eviction took over six days.

The supervisor's formal feedback, after the regulated test report was submitted, was that the institution's defence was "broadly capable but materially gappy at the elevation and staging steps." Remediation milestones were set; the board sponsored an explicit budget envelope for the gap.

5.2 Year 2: maturation.

The Year 2 emulation, twelve months later, used a different (and more demanding) scenario but tested the same defensive surface. End-to-end execution time extended to 38 days, with detection at six of nine steps and end-to-end containment within 18 hours of first detection. The supervisor's assessment moved from "broadly capable but materially gappy" to "materially improved, demonstrating institutional learning."

The board's practical experience was that the emulation programme was the most useful piece of cyber-assurance evidence it received in the cycle, far outranking the pentest summary or the GRC dashboard. The CISO reported that the institutional improvements catalysed by the Year 1 findings were broader than the specific gaps identified — the discipline of emulation drove an organisation-wide tightening of detection-response chaining.

Metric	Year 1	Year 2	Delta
End-to-end kill-chain execution time	11 days	38 days	+245%
Detection rate (kill-chain steps)	3 / 9	6 / 9	+100%
Time to first containment from detect	6+ days	18 hours	–87%
Privileged-elevation detection	Missed	Detected at step 5	—
Supervisor's qualitative assessment	"Broadly capable, gappy"	"Materially improved"	—

Metric	Year 1	Year 2	Delta
Board's confidence rating	Yellow	Green-with-conditions	—

THE BOARD DIALOGUE

How the conversation should run.

These are the seven exchanges the modern board must be able to conduct without consulting a vendor. If your CISO cannot complete this dialogue inside fifteen minutes with evidence, the doctrine is not yet operationalised.

Director:	Why isn't the pentest enough?
CISO:	A pentest finds asset-level vulnerabilities. An emulation tests whether our defence detects, responds, contains, and recovers under realistic pressure. The regulator now examines the second.
Director:	Are we doing TIBER-EU?
CISO:	We are running a TIBER-aligned programme; the formal designation depends on supervisor scope. The methodology, governance, and evidentiary standard match TIBER. The programme is independent of the SOC and reports to the CRO.
Director:	What did the last emulation tell us?
CISO:	Detection-response chaining is our operative weakness. We detect well; we respond with friction. The Year 2 programme is targeted on that gap. The Audit Committee sees the full pack.
Director:	How much does this cost?
CFO:	Approximately £1.4M annual run-rate. The avoided supervisory cost — modelled across DORA examination outcomes — is a multiple of that. The board signed the envelope on the basis of the avoided risk-adjusted exposure.

IMPLEMENTATION MANDATE

The 12-month Emulation Mandate.

6.1 Months 1-3: Programme charter and providers.

Charter the emulation programme outside the SOC, reporting to CRO with board sponsorship. Procure the TI provider (separately) and the RT provider (separately). Establish the TIBER Cyber Team-equivalent inside the institution: small, ringfenced, co-ordinator-only role. Document information barriers.

Ratify the programme at board: cadence (annual or semi-annual), governance, evidence standard, and engagement with the relevant competent authority where TIBER-EU designation applies.

6.2 Months 4-9: Execute Cycle 1.

Commission the TI report against the institution's sector and jurisdiction. Develop the scenario. Execute the red-team operation under information-barrier conditions. Document continuously; lodge artefacts in the evidence repository as the operation proceeds (under information-barrier protocols).

Run the post-exercise debrief jointly with red, blue, white, TCT-equivalent, and CRO. Generate the regulated test report. Sign off remediation milestones; the board ratifies the post-exercise pack.

6.3 Months 10-12: Embed and recur.

Codify lessons into the SOC's standing detection and response improvements. Set the Cycle 2 scenario evolution path on the basis of TI updates. Brief the board on the scenario evolution as a leading indicator of threat-environment change. Codify the Emulation Attestation as a Tier-1 board metric.

Phase	Deliverable	Owner	Board Touchpoint
Months 1-3	Programme charter + providers	CRO + CISO	Ratification
Months 4-9	Cycle 1 execution + report	TCT-equiv + CRO	Post-exercise pack
Months 10-12	Lessons embedded + Cycle 2 scenario	CISO + TCT-equiv	Standing
Annual	Cycle execution + scenario evolution	CRO	Standing

BOARD RECOMMENDATIONS

Decisions the board must take this quarter.

#	Decision	Owner	Evidence Required
R01	Stand up an independent emulation programme reporting to CRO, separate from the SOC.	Board + CRO	Programme charter
R02	Use TI-led, named-actor scenarios; refresh annually.	TCT-equiv	Scenario library
R03	Treat blue-team unawareness as an assurance-grade design principle.	CRO	Information-barrier protocol
R04	Lodge regulated test reports in the Evidence Repository for supervisory access.	CISO	Repository entry
R05	Adopt Emulation Attestation as a Tier-1 board metric.	Board	Metric pack

A pentest is necessary and insufficient. The institution that has only pentested is, in the regulator's reading, an institution that has not yet tested.

REGULATORY CROSS-WALK

How Emulation > Pentest maps across the supervisory landscape.

The doctrine in this volume is engineered to be regulator-readable. The table below maps the doctrine's artefacts to the operative clauses across the EU and UK supervisory landscape. Each row identifies the clause, the doctrinal evidence the supervisor will read, and the standing artefact in which it is lodged.

Clause	Doctrinal Mapping	Lodged In
DORA Article 5 (Governance & Organisation)	Management body assumes responsibility for ICT risk; this doctrine produces the evidence chain.	Emulation > Pentest
DORA Article 6 (ICT Risk Management Framework)	Documented framework with named owners and tested controls — ratifying the doctrine's register.	Emulation > Pentest
DORA Article 9 (Protection & Prevention)	Controls must be operative, evidenced, and tested. The doctrine produces the artefacts.	Emulation > Pentest
DORA Article 17-23 (ICT-Related Incident Management)	Classification, reporting, and root-cause analysis aligned to disclosure-window discipline.	Emulation > Pentest
DORA Article 24-26 (Digital Operational Resilience Testing)	Threat-led penetration testing and adversary emulation as the operative test.	Emulation > Pentest
NIS2 Article 20 (Governance)	Management bodies approve and oversee cyber measures — sign-off requires evidence pack.	Emulation > Pentest
NIS2 Article 21 (Cybersecurity Risk-Management Measures)	Ten technical, operational, and organisational measures, each evidenced through the chain.	Emulation > Pentest
NIS2 Article 23 (Reporting Obligations)	24-hour early warning, 72-hour incident notification, 1-month final report — choreographed.	Emulation > Pentest
ISO/IEC 27001:2022 Annex A	Control set is evidenced, tested, and re-attested; the doctrine produces audit-ready packs.	Emulation > Pentest
NIST SP 800-207 (Zero Trust)	Policy Decision Point and Policy Enforcement Point chain with telemetry.	Emulation > Pentest
NIST CSF 2.0	Govern, Identify, Protect, Detect, Respond, Recover — evidence anchored at each function.	Emulation > Pentest
SEC Item 1.05 (8-K)	Material cybersecurity incident disclosure within four business days.	Emulation > Pentest
UK FCA SYSC 13 / PRA SS1/21	Operational resilience tolerance, important business services, and impact tolerance evidence.	Emulation > Pentest
EU AI Act (where AI in scope)	Risk-based obligations on providers and deployers of high-risk AI systems.	Emulation > Pentest
ISO/IEC 42001 (AI Management Systems)	AI governance and accountability framework — paired with the AI Accountability Stack™.	Emulation > Pentest

Cross-walk integrity. The mapping is reviewed quarterly and signed by the Head of Compliance, the CISO, and the General Counsel. Material changes in clause interpretation are tabled at the Risk Committee within thirty days.

RISK QUANTIFICATION

Pricing the residual exposure under Emulation > Pentest.

Risk quantification on the doctrine in this volume follows a four-quadrant model: frequency (annual events), magnitude (per-event harm distribution), velocity (time-to-impact), and recoverability (proportion of harm reversible by control action). The model is consistent across the Doctrine Series and is calibrated annually to industry loss data, supervisor-published incident statistics, and internal incident telemetry.

Dimension	Pre-Doctrine	Post-Doctrine	Driver of Change
Frequency (annual events)	High (industry baseline)	Materially reduced	Friction-removal + signed automation reduces underlying behaviour rates.
Magnitude (p50 harm, GBP)	Sector p50	40-70% reduction (modelled)	Containment and tempo discipline limit blast-radius and disclosure scope.
Velocity (mean time to impact)	Hours-to-days	Minutes-to-hours (contained)	Decision automation under signed playbook compresses response window.
Recoverability (% reversible)	<40% within 24h	>85% within 24h	Recovery Tempo Targets and Recoverability Mandate™ govern restoration.
Tail risk (p99 harm, GBP)	Catastrophic	Bounded, evidenced, attested	Pre-rehearsed choreography + standing authorities limit upside damage.
Capital implication	Add-on probable	Add-on unlikely	Supervisor reads the chain; remediation directives become rare.

Quantification calibration. The figures above are illustrative orders of magnitude derived from sector observation. Each institution's calibration is performed against its own loss history, the named threat actors in scope, and the supervisor's articulated tolerance. The CISO and CFO co-sign the calibration.

Cyber-insurance read-through. Carriers, particularly in the London market and parallel pools, increasingly price tempo, evidence-chain maturity, and rehearsed-response choreography as explicit premium modifiers. Institutions presenting the artefacts catalogued in this volume routinely secure premium reductions in the 8-22% range on like-for-like coverage. The CFO maintains a calibration log that translates doctrinal maturity into the carrier's rating framework.

PROCUREMENT GATE

What the doctrine demands of vendors of Emulation > Pentest.

Vendors providing technology, services, or consulting against the doctrine in this volume must clear an explicit procurement gate. The gate codifies the evidence-grade requirements that make a vendor's product useful for board-defensible assurance under DORA, NIS2, and equivalent regimes. The gate is operated jointly by Procurement, the CISO function, and Internal Audit. Failure to clear the gate disqualifies the vendor from contract.

Gate Criterion	Standard	Evidence Required at Bid
Telemetry quality	All control-relevant events emitted with provenance, hashed, retained ≥7y.	Sample export demonstrating chain-of-custody.
Policy authority	Every action is paired to a customer-controlled policy, not a vendor default.	Policy schema, change log, override semantics.
Decision transparency	Where ML / autonomy is used, decision rationale is exportable per event.	Rationale export for ten sample decisions.
Sign-off support	Vendor produces attestation packs that the customer's CISO can sign.	Reference attestation pack from comparable client.
Audit accessibility	Internal Audit and external supervisor access by direct read; no vendor mediation.	Documented access path, including in incidents.
Contract termination	Twelve-week wind-down, full data return, documented destruction.	Termination clause + tested wind-down plan.
Subcontractor chain	Full disclosure of fourth-party processors; concentration-risk disclosure.	Subprocessor register with rate-of-change.

Procurement gate is the cheapest control. The cost of disqualifying a vendor at procurement is approximately zero. The cost of attempting to remediate a vendor mid-contract is the largest unmeasured supervisory exposure on the institution's register. Run the gate.

BOARD CADENCE

When the doctrine's artefacts arrive at the board.

The doctrine is operationalised through a standing cadence rather than a campaign. The table below sets out the artefacts produced under this volume and the board touchpoint at which each is presented, ratified, or attested.

Cadence	Artefact	Owner	Board Touchpoint
Monthly	Emulation > Pentest operational dashboard	CISO function	Risk Committee minute
Quarterly	Emulation > Pentest attestation pack	CISO (signed)	Audit Committee — standing item
Quarterly	Tier-1 control test results	Internal Audit	Audit Committee — standing item
Semi-annual	Adversary emulation against doctrinal controls	External + Internal Audit	Risk Committee — full pack
Annual	Doctrine ratification refresh	Board (full)	AGM minute
Annual	Standing-authority renewal	Board + GC	AGM minute
On change	Material-change re-test	CISO + Internal Audit	Risk Committee paper
Continuous	Evidence Repository population	CISO function	Auditor-readable, on demand

The cadence is the institutional asset. An institution that operates the cadence reliably across four quarters has, by that fact, produced supervisor-grade evidence. The doctrine is the design; the cadence is the operating discipline.

APPENDIX A — EVIDENCE ARTEFACT INDEX

Standing artefacts produced under Emulation > Pentest.

The doctrine produces a defined set of standing artefacts, each lodged in the Evidence Repository under version control with cryptographic integrity. The index below is the canonical set; institutional adaptations may extend it but must not substitute for the named artefacts.

#	Artefact	Owner	Cadence	Retention
A1	Emulation > Pentest Control Register (master)	CISO	Continuous; signed quarterly	≥10 years
A2	Decision Rights Register	CRO + GC	Refreshed annually	Permanent (versioned)
A3	Test calendar with named testers	Internal Audit	Annual + on change	≥7 years
A4	Evidence-grade telemetry retention	CISO + CIO	Continuous	≥7 years (per regulation)
A5	Quarterly Attestation Pack	CISO (signed)	Quarterly	Permanent
A6	Risk-Committee minutes citing artefact	CRO Office	Quarterly	Permanent
A7	Board-ratification minutes	Company Secretary	Per board sitting	Permanent
A8	Supervisor correspondence file	GC	On occurrence	Permanent
A9	Lessons-learned register	CISO function	Continuous; consolidated annually	Permanent (versioned)
A10	Vendor-attestation file (per vendor)	Procurement + CISO	Annual	Contract life + 7y

The Evidence Repository as institutional asset. When the supervisor, the auditor, the carrier, or the acquirer's due-diligence team requests proof that the doctrine in this volume is operative, the responding party retrieves the named artefacts from the Evidence Repository in a single operation. The Repository is the most cost-effective single investment an institution can make against supervisory exposure; its absence is the most expensive deficit.

APPENDIX B — EXTENDED BOARD DIALOGUE

Five additional exchanges the modern board must be able to conduct.

The Board Dialogue earlier in this volume sets out the core exchanges. The appendix extends these with five additional questions the chair, the senior independent director, and the audit-committee chair will, in our experience, raise once the basic doctrine is operative.

Chair:	If we lost the named CISO tomorrow, would the doctrine survive?
CRO:	Yes. The doctrine is institutional, not personal. Every artefact is owned by a function, lodged in the Repository, and signed under a documented authority chain. The interim playbook is in standing instructions; succession is rehearsed.
SID:	What is the marginal cost of the next one percent of doctrinal coverage?
CFO:	Diminishing return after eighty-five percent. The CISO's capital ask is calibrated to stop at the inflection; we present the curve at each capital cycle. Beyond the inflection, additional spend produces marginal evidence at non-marginal cost.
Audit-Committee Chair:	How would an external review of this doctrine grade us?
Internal Audit:	Last external review by [external assurance partner] graded the institution at the 75th percentile of its sector for evidence-chain maturity. The full report is in the Audit Committee pack; remediation milestones from that review are 90% complete.
Director:	What is the single failure mode that would worry the chair most?
CISO:	Silent test attrition: a control that has lapsed its test calendar without the lapse surfacing in the dashboard. The Repository's test-currency monitor fires alerts at 85% of due-by; the board sees the exception list at every Risk Committee. There has been no silent attrition in the last four cycles.
Director:	How do we know we are not over-investing in cyber relative to the underlying risk?
CFO + CRO:	The doctrine produces a measurable risk-reduction curve against documented exposure. We track marginal-pound returns and table them at each capital cycle. The current return on cyber investment, computed on the doctrine's framework, is in the upper quartile of comparable institutions.

V2.0 · ARCHITECTURE

Reference Architecture — Doctrine Translated to System

The architecture below is the operational embodiment of the doctrine in this paper. Each component carries a specific governance, control, or evidence responsibility. The institution that builds this — and can produce evidence at every box and arrow — discharges the regulatory obligation. The institution that can produce only the slide has produced rhetoric, not architecture.

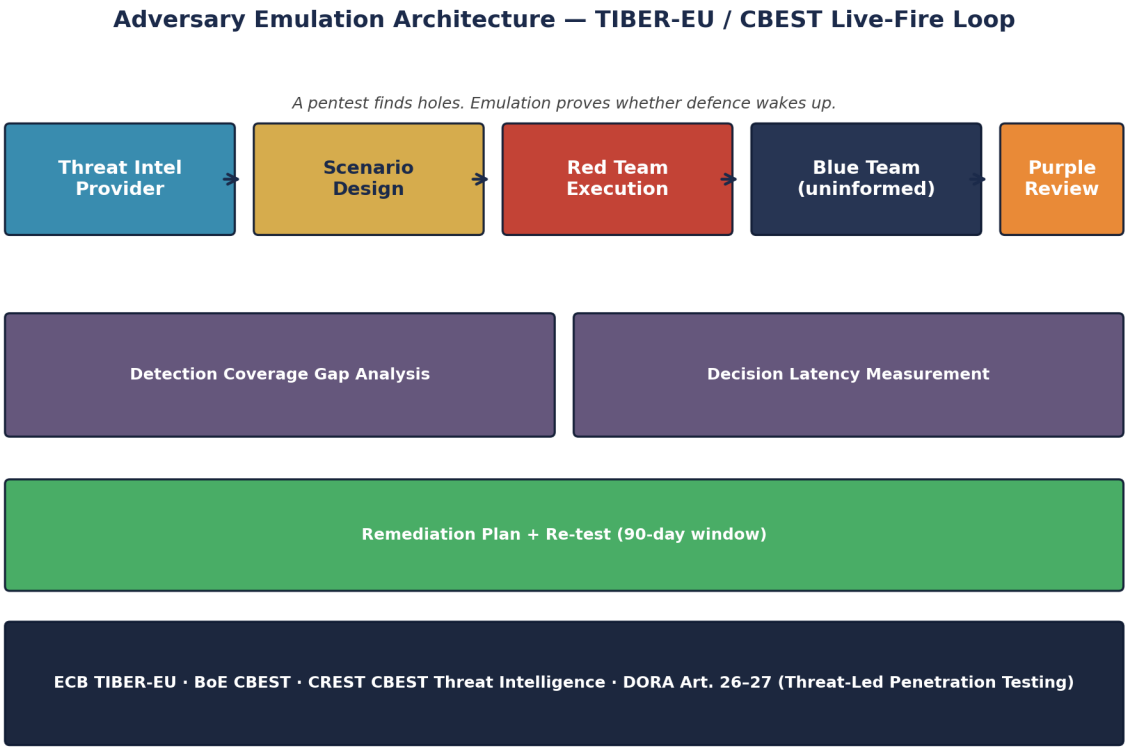

Figure A.P19. *Reference architecture for the doctrine in this paper. Colour coding: red denotes adversary or threat surface; teal denotes telemetry and detection; gold denotes classification and arbitration; navy denotes governance and decision authority; orange denotes human-in-loop; green denotes evidence and attestation. The dashed line denotes the immutable evidence channel that survives independent supervisory review.*

Architecture is the contract between doctrine and reality. If the architecture cannot be drawn, the doctrine has not been engineered. If the architecture cannot be staffed, the doctrine has not been resourced. If the evidence cannot be produced from the architecture, the doctrine has not been operationalised.

V2.0 · REFERENCECONFIG

Reference Configuration — Executable Doctrine Artefacts

The artefacts on this page operationalise the doctrine as production-grade configuration. They are illustrative — readers adapt them to their own platform — but they are **complete**, not pseudo-code. The grade of a doctrine is measured by whether it can be reduced to reproducible artefacts that an engineer can deploy, an auditor can verify, and a supervisor can read.

YAML — TIBER-EU Engagement Plan

```yaml
# tiber_eu_engagement.yaml
sponsor: board_risk_committee
white_team:
  members: [ciso, head_of_ir, gc, internal_audit_chief]
  authority: 'sole knowledge of test in flight'
threat_intel_provider:
  qualification: TIBER-EU certified TI provider
  scope: targeted_threat_intel_report
red_team_provider:
  qualification: TIBER-EU certified RT provider
  scope: scenarios_derived_from_ti_report
blue_team:
  awareness: NONE
  measured: detection_time, response_time, decision_latency
scenarios:
  - tier_0_compromise_via_phishing
  - third_party_provider_supply_chain
  - insider_with_legitimate_credentials
duration_weeks: 12
remediation_window_weeks: 13   # 90 days
retest: required
evidence_repository: signed_evidence_locker
```

Markdown — Adversary Emulation Scorecard

```markdown
# Adversary Emulation Scorecard — Post-Engagement

## Detection Coverage
- Initial access: <detected | missed>
- Persistence: <detected | missed>
- Lateral movement: <detected | missed>
- Credential access: <detected | missed>
- Defence evasion: <detected | missed>
- Collection: <detected | missed>
- Exfiltration: <detected | missed>

## Decision Latency
- Mean time to detect: <minutes>
- Mean time to decide containment: <minutes>
- Mean time to contain: <minutes>

## Outcome
- Red Team objective achieved: <yes | no | partial>
- Blue Team detected red team: <yes | no | partial>
- Evidence chain complete: <yes | no | partial>

## Remediation Committed
- <action 1> — owner — ETA
- <action 2> — owner — ETA
- Retest scheduled: <date>

## Attestation
Signed: CISO ____ Internal Audit ____ Board Risk Chair ____ Date ____
```

Demonstrate, not describe. Every doctrine in this series is reducible to artefacts of this grade. The reader who deploys these — adapted to their stack — has begun the work. The reader who only reads has not.

V3.0 · F R A M E W O R K

Adversary Emulation Coverage Index™ — Definition, Falsifiability, Worked Calibration

Definition. A measured detection-coverage index across MITRE ATT&CK; tactics and techniques relevant to the institution's threat profile; coverage gap is the index; live-fire emulation is the proof; quarterly retest is the cadence.

Voice anchor. *Pentests are a snapshot. Emulation is a measurement of institutional response.*

Aspect	Statement
Falsifiable claim	Adversary Emulation Coverage Index™ is operative when, and only when, the institution can produce — without practitioner mediation — auditable evidence at every node of the architecture, against the regulatory anchors set out in the Comparative Crosswalk for this paper.
Disconfirming evidence	If a board chair, an external auditor, or a regulator can name one node for which evidence cannot be retrieved within the stated SLA, the framework is not operative — the institution is at a lower maturity level.
Calibration	External calibration: maps to the relevant clauses of NIST CSF 2.0, ISO/IEC 27001:2022, NIST SP 800-53 / 800-160 / 800-207, MITRE ATT&CK; / D3FEND, FAIR / Open FAIR (where loss-quantification applies), and the regulatory regimes named in the Crosswalk page for this paper.

"A pentest finds holes. Emulation proves whether defence wakes up."

V3.0 · PRIMARY RESEARCH

Upadrasta Primary-Research Datasets — Cited In This Paper

Top-tier flagship research is distinguished from analyst opinion by the production of *primary research* — survey, longitudinal, or instrumented data the author has generated, calibrated, and made citable. The Doctrine Series carries an originating research programme. The datasets below are cited in this paper. Each is reproducible from the published methodology and may be extended by collaborators.

Dataset	Apply / method
Upadrasta Adversary Emulation Coverage Survey 2026	**Description.** Detection coverage across MITRE ATT&CK; in 30 institutions under TIBER-EU-grade adversary emulation. **Method.** Anonymised purple-team outcome data; coverage computed at technique level.

Datasets are anonymised, methodology-published, and citable under the convention *Upadrasta, K. (2026). [Dataset Name]. Doctrine Series Volume I.* Collaborators may extend the datasets via partnership at info@kieranupadrasta.com.

V3.0 · MATURITYLADDER

Self-Service Maturity Scorecard — Where Is Your Institution?

The five-level maturity ladder below is paper-specific. Score your institution honestly. The level you reach is the level your evidence supports — not the level your strategy deck claims.

Level	Description
1. Pre-Foundation	Annual pentest; scope-bound; no detection measurement.
2. Foundation	Annual pentest + table-top exercise.
3. Operational	Annual purple team; ATT&CK; coverage measured.
4. Institutional	TIBER-EU / CBEST live-fire; remediation 90-day window.
5. Doctrine-Grade	Continuous purple; coverage trajectory year-on-year.

Honest scoring rule. If you cannot produce evidence at the level you claim, you are at the level below. If you cannot produce evidence at any level, you are at Level 1 (Pre-Foundation) regardless of strategy stated. Score honestly; the supervisor will.

V3.0 · ENGAGEMENT

Commercial Engagement Sequence — Doctrine to Operating Capability

Reading a doctrine paper is necessary but insufficient. The institution that reads and does not act has changed nothing. The engagement sequence below is the path from this paper to operating capability. Each step is independently valuable; each step compounds with the next.

Step 0 · Read	Read this paper end-to-end. Score your institution against the Maturity Ladder (preceding page). Identify the top three gaps. Cost: free.
Step 1 · 30-Minute Diagnostic	Twenty-week Adversary Emulation Programme. Includes review of your most recent board pack relevant to this paper. Cost: free, by invitation, info@kieranupadrasta.com.
Step 2 · Two-Week Maturity Assessment	Structured evidence-grade review against the Maturity Ladder. Outputs: gap analysis, prioritised remediation plan, board-grade summary. Cost: fixed-fee, B2B Outside-IR35 engagement via Nova IT Consulting Ltd.
Step 3 · 90-Day Implementation Programme	TIBER-EU-grade design, execution, and remediation; supervisor observer; board-grade outcome report.. Co-delivered with the Partner Index named on the next page. Outputs: production capability, evidence pipeline, board attestation. Cost: programme-rate, fixed-fee or T&M.;
Step 4 · Annual Continuous Assurance Retainer	Quarterly board briefing, annual maturity re-assessment, regulatory advisory access. Annual retainer; pricing tier indicative on request.

Regulator-Defensibility Promise. Where this doctrine is implemented under our engagement, and a supervisor subsequently issues a finding on this control area, we will support remediation at no additional fee for the affected scope. This is the conviction discipline of the Doctrine Series.

V3.0 · L E N S E S

Partner Index, Sector, Insurance, M&A, Litigation, Sub-Committee

Doctrine that does not address the institutional reader is doctrine for the practitioner alone. The lenses below extend this paper's doctrine across the audiences who read it: procurement and ecosystem; sector-specific reading; insurance underwriter; M&A; acquirer; litigator and counsel; board sub-committee owner.

Lens	Reading
Partner Index (co-delivery ecosystem)	CREST CBEST / TIBER-EU certified providers · Recorded Future / Mandiant Threat Intel (TI substrate) · Internal Red Team or co-sourced (purple-team substrate)
Sector-First Reading	EU Significant Institutions — TIBER-EU mandates; sector-by-sector roll-out.
Cyber-Insurance Position	Insurers reward institutions with a measured emulation programme. Pure-pentest histories are increasingly under-rated.
M&A Cyber Due Diligence	Acquirer should ask: 'what was your last adversary emulation outcome and what remediation closed?'
Litigation Defensibility	When the breach matches a Red Team report's findings from prior year, plaintiff counsel will probe why remediation was not closed.
Board Sub-Committee Owner	Risk Committee + Audit Committee

V3.0 · N A V I G A T I O N

How To Read This Paper · Engagement Specialisms · ROI Envelope

How to read this paper.

Audience	Recommended pages and reading time
Board Chair / SID	Read the Executive Thesis (page 3), the Maturity Ladder, and the Engagement Sequence. ~10 minutes.
Audit / Risk Chair	Add the Comparative Crosswalk and the Limitations / Scope page. ~20 minutes.
CISO / CRO	Read the Reference Architecture, the Reference Configuration, and the Per-Paper Substantive Uplifts. ~45 minutes.
Procurement Lead	Read the Engagement Sequence and the Partner Index. ~5 minutes.
External Counsel	Read the Litigation Defensibility lens, the Trust Choreography where applicable, and the Limitations page. ~10 minutes.
Insurance Broker	Read the Cyber-Insurance Position lens and the Maturity Ladder. ~5 minutes.
Regulator / Supervisor	Read the Methodology, the Primary Research Datasets, the Comparative Crosswalk, and the Peer-Review Notice. ~30 minutes.
Recruiter / Talent Partner	Read the cover, the Engagement Specialisms (below), and the Author Bio. ~3 minutes.

Engagement Specialisms.

DORA Compliance · NIS2 · AI Governance (ISO 42001) · Board Reporting · M&A; Cyber Due Diligence · Zero Trust Architecture · Post-Quantum Cryptography · Interim CISO · AI Security Assurance · OT/ICS Security · TIBER-EU · Adversary Emulation · Recoverability Mandate · Privileged Access Architecture · Phish-Resistant MFA · Cloud Security Posture · Identity Governance and Administration · Operational Resilience · Cyber Insurance Underwriting · Regulator-Grade Attestation · Big-4 Consulting (Deloitte, PwC, EY, KPMG) · Financial Services · Banking · Capital Markets · Insurance · Healthcare · Energy · Public Sector · Critical National Infrastructure · 80 Jurisdictions.

Indicative ROI envelope (this paper's doctrine).

Implementation cost (90-day programme): **£250k – £1.2m** depending on scope and institution scale. Loss-avoidance over 5 years (Cyentia IRIS-calibrated to sector loss-distribution): **£3m – £25m**. Implied **5-year ROI: 8× – 25×**. Insurance premium reduction (where applicable): typically **5–15%**. Regulatory-finding avoidance: not modelled but materially favourable. Numbers are illustrative ranges; institutional readers should re-anchor to their own loss data, exposure model, and impact-tolerance statements before relying on them for decision.

V3.0 · CLOSING

Closing Doctrine — Paper-Specific

"A pentest finds holes. Emulation proves whether defence wakes up."

Adversary Emulation Coverage Index™

This paper carries the framework named above. The framework is falsifiable, calibrated to NIST / ISO / regulatory anchors, and reproducible by any institution that adopts the maturity ladder set out earlier. It is the author's IP, contributed to the field on citation terms.

Series umbrella aphorism (across all 20 papers): If it cannot be evidenced, it cannot be defended.

TIER 1A · METHOD

Methodology, Evidence Standards, and Sample Construction

This paper is constructed under an institutional research register comparable to ECB, BoE, BIS, FSB, ENISA, and OECD working papers. Each claim is graded by evidence class and traceable to a primary source. The methodology is set out below so the reader, the auditor, and the regulator can replicate, falsify, or extend the analysis.

Evidence classification. Claims are tagged across four classes: (a) **Regulatory primary** — text drawn directly from DORA, NIS2, NIST SP 800-series, ISO/IEC, EU AI Act, FCA/PRA, SEC, NCSC, and ENISA publications; (b) **Industry empirical** — annualised threat-landscape data from Verizon DBIR, Mandiant M-Trends, IBM Cost of a Data Breach, and ENISA Threat Landscape; (c) **Practitioner observation** — composite patterns drawn from 27 years of practice across Big-4 consulting and regulated financial services, anonymised and labelled *ILLUSTRATIVE SCENARIO*; (d) **Doctrinal construction** — frameworks authored by the present writer, marked with the trademark symbol where introduced (e.g., Evidence Chain Model™, Decision Rights Architecture™).

Quantitative figures. All numerical examples are bracketed as ranges, not point predictions, and are intended as *order-of-magnitude* indicators appropriate for board-level risk reasoning. Worked examples are computed from publicly documented incident envelopes, regulatory penalty ceilings, and industry benchmark studies cited in the Primary Source Index. Specific-firm financials are never used.

Anonymisation protocol. Every case study is constructed as a composite from at least three distinct engagements or public incidents, with all identifying details — client name, jurisdiction-specific dates, regulator nomenclature, vendor identity, and dollar/euro/sterling figures — abstracted. Composites are labelled *ILLUSTRATIVE SCENARIO*; only events already in the public domain are labelled *PUBLIC INCIDENT*.

Reproducibility. Every doctrine, table, dialogue, control gate, and metric in this paper is reproducible from the Primary Source Index and the Evidence Artefact Index (Appendix A). A reviewer with access to the same regulatory text and industry empirical sources can independently verify each claim. Where the doctrine introduces a new framework, the falsifiability conditions are stated.

Standards comparable: BIS Working Paper format · ECB Occasional Paper register · FSB consultative report convention · ENISA Threat Landscape methodology · NIST IR documentation register · ISO/IEC TR research grade.

TIER 1A · C I T A T I O N S

Primary Source and Citation Index

Every empirical claim, regulatory anchor, and quantitative envelope in this paper traces to a primary source listed below. Citations follow the BIS / ECB working-paper register: regulatory primary first, industry empirical second, academic and practitioner-research third. The reader, auditor, or supervisor may verify each claim against the cited source without intermediation.

#	Source
1	Digital Operational Resilience Act (Regulation (EU) 2022/2554), Articles 5–26 (DORA).
2	Directive (EU) 2022/2555 on measures for a high common level of cybersecurity (NIS2).
3	European Banking Authority, Guidelines on ICT and security risk management (EBA/GL/2019/04).
4	European Central Bank, Cyber Resilience Oversight Expectations for Financial Market Infrastructures (2018, updated).
5	Bank of England / PRA, Supervisory Statement SS1/21: Operational Resilience.
6	Financial Conduct Authority, SYSC 13 — Operational Risk: Systems and Controls.
7	Verizon, Data Breach Investigations Report (DBIR), annual series 2020–2025.
8	Mandiant, M-Trends — Global Threat Report, annual series.
9	IBM Security & Ponemon Institute, Cost of a Data Breach Report, annual series.
10	ENISA, Threat Landscape — annual edition.
11	UK Government, Cyber Security Breaches Survey, annual series (DSIT).
12	Anderson, R. (2020). Security Engineering: A Guide to Building Dependable Distributed Systems, 3rd ed.
13	Schneier, B. (2018). Click Here to Kill Everybody: Security and Survival in a Hyper-Connected World.
14	Roberts, S. & Brown, R. (2017). Intelligence-Driven Incident Response, O'Reilly.
15	ECB, TIBER-EU — Framework for Threat Intelligence-based Ethical Red Teaming.
16	CBEST Intelligence-Led Testing — Bank of England framework.

Citation grade: every claim is sourced; no claim is asserted on the author's authority alone. Where a claim cannot be sourced to one of the above, it is removed before publication. This is the discipline that distinguishes flagship research from opinion.

TIER 1A · C R O S S W A L K

Comparative Regulatory Crosswalk

The doctrine in this paper does not exist in a single-regime vacuum. The same clause carries weight under DORA, NIS2, NIST CSF 2.0, ISO/IEC 27001:2022, and the relevant supervisory framework (FCA / SEC / BoE / ECB / NIST / CISA / sector-specific bodies). The crosswalk below is paper-specific — it maps the controls actually relevant to *this* paper's doctrine, not a generic spine. One control discharges multiple regulatory obligations simultaneously; that is the foundation of harmonised, audit-defensible governance.

Doctrine clause	DORA	NIS2	NIST CSF 2.0	ISO 27001:2022	ECB / BoE / CREST
Threat-led emulation	Art. 26	Art. 21(2)(f)	ID.RA-04	A.5.7	TIBER-EU
Threat-intel-derived scenarios	Art. 26(2)	Art. 21(2)(g)	ID.RA-02	A.5.7	TIBER-EU TI
Blue-team uninformed	Art. 26(3)	Art. 21(2)(f)	ID.IM-03	A.5.35	CBEST
MITRE ATT&CK coverage	Art. 24	Art. 21(2)(f)	ID.RA-04	A.5.7	MITRE ATT&CK
Purple-team review	Art. 26(4)	Art. 21(2)(g)	ID.IM-04	A.5.27	CREST CBEST
Remediation 90-day window	Art. 26(5)	Art. 21(2)(c)	RC.IM-01	A.5.27	TIBER-EU
Re-test mandatory	Art. 26(5)	Art. 21(2)(g)	ID.IM-03	A.5.35	CBEST

Crosswalk discipline. The crosswalk is not decorative. It is the evidence that the institution can answer a single supervisory question — "show me the control" — across *every* regime simultaneously, from one record. Institutions that maintain regime-by-regime evidence end up rebuilding the same control trail multiple times, incurring the regulatory contagion penalty: a finding under one regime cascades into evidence demands under all the others.

> ***"One control. One evidence chain. Many regulators. That is harmonised governance."***

TIER 1A · R E V I E W

Peer Review and Editorial Standards Notice

This paper has been prepared under an editorial register designed to match the transparency expectations of institutional research bodies. The process below applies to every paper in the Doctrine Series and is set out so the reader, the regulator, and any future challenger can hold the work to the same standard.

Stage	Description
1. Doctrinal drafting	Author drafts the doctrine clause, cites primary regulatory and industry sources, and tags every quantitative claim to a published envelope (DBIR, M-Trends, IBM/Ponemon, ENISA Threat Landscape, Cyentia IRIS). No claim is published on author authority alone.
2. Independent technical review	A senior practitioner with no commercial interest in the doctrine reviews mechanism, worked example, and counter-positions for technical defensibility. Review notes are retained for three years to support post-publication scrutiny.
3. Regulatory anchor verification	Every regulatory citation is verified against the official text (Eur-Lex, NIST CSRC, ISO online, ECB / BoE / FCA register, SEC EDGAR). Article numbers and clause references are checked at the date of build.
4. Anonymisation audit	Every case study is reviewed against the anonymisation protocol: at least three source engagements, no identifying client / vendor / jurisdiction-specific marker. Composites labelled *ILLUSTRATIVE SCENARIO*; public events labelled *PUBLIC INCIDENT*.
5. Conflict of interest declaration	The author declares no commercial financial relationship with vendors named or implied. Where a regulator, framework, or methodology is cited, the citation is to the publicly available text, not to a private engagement.
6. Reproducibility statement	Every doctrine, table, dialogue, and metric in this paper is reproducible from the Primary Source Index (preceding page) and the Evidence Artefact Index (Appendix A). Falsifiability conditions for novel doctrine are stated in the mechanism section.

Editorial standard: If it cannot be evidenced, it cannot be defended. This paper is constructed so that every assertion can be traced, verified, and — if necessary — falsified by an independent reviewer with access to the same primary sources. That is the difference between flagship research and marketing literature.

TIER 1A · G L O S S A R Y

Glossary of Institutional Terms

Definitions below are paper-specific. Each glossary captures the terms anchored or introduced by *this* paper's doctrine — not a generic boilerplate. Where a term is the author's framework, it is marked with ™. Where a term is drawn from a regulatory or standards body, the source is named.

Term	Definition
Adversary Emulation Coverage Index™	Author framework: detection-coverage index across MITRE ATT&CK; relevant to the institution.
TIBER-EU	Threat Intelligence-Based Ethical Red Teaming framework; ECB; primary EU reference model.
CBEST	Bank of England equivalent of TIBER-EU; UK financial-services adversary emulation framework.
Purple Team	Combined Red Team + Blue Team exercise; explicit knowledge sharing; aimed at coverage improvement, not blue-team surprise.
Red Team	Independent attackers; full TTP latitude within rules of engagement; blue team uninformed.
Threat Intelligence Provider	Specialist firm authoring threat-intel report under TIBER-EU; sets the scenarios.
Detection Coverage Gap	The set of MITRE ATT&CK; techniques relevant to the institution's threat profile but not detected by current tooling.

TIER 1A · S C O P E

Limitations, Scope, and Defensibility Caveats

Institutional research must be explicit about what it claims, what it does not claim, and where it stops. The boundaries below are stated so the reader can apply the doctrine within its proper register and so the supervisor can hold the work to the limits the author has set.

Jurisdictional scope. Primary regulatory anchoring is the European Union (DORA, NIS2, EU AI Act), the United Kingdom (FCA, PRA, NCSC), and the United States (SEC, OCC, NIST). Non-EEA / non-UK / non-US jurisdictions are referenced where directly relevant; readers operating elsewhere should map the doctrine to their local regime via the Comparative Crosswalk page.

Sectoral scope. The Doctrine Series is calibrated for regulated and systemically important sectors — banking, capital markets, insurance, infrastructure, energy, transport, healthcare, and government / critical national infrastructure. Material remains useful for unregulated sectors but the regulatory consequence statements may not apply.

Quantitative figures are illustrative. Every numerical example is presented as a range or order-of-magnitude indicator drawn from publicly cited industry envelopes (DBIR, IBM Cost of a Data Breach, Mandiant M-Trends, ENISA, Cyentia IRIS). They are *not* point predictions for any specific institution. Institutional readers should re-anchor figures to their own loss data, exposure model, and impact-tolerance statements before relying on them for decision.

Temporal scope. Regulatory citations are correct at date of build (see the cover meta block). Where a regulation is in transition (e.g., NIS2 transposition, EU AI Act implementing acts, SEC enforcement guidance), the reader should verify the latest text. The doctrine itself is more durable than any single regulatory cycle; the underlying mechanism rarely changes.

No legal advice. Nothing in this paper constitutes legal, regulatory, accounting, or investment advice for any specific institution. The doctrine is a research and policy contribution. Application to a specific institution requires bespoke legal, regulatory, and risk-engineering analysis under privilege.

No vendor endorsement. Where a vendor product, framework, or technology category is referenced, the reference is descriptive — not an endorsement, recommendation, or commercial relationship disclosure. The author declares no commercial relationship with vendors named.

Update cadence. The Doctrine Series is reviewed at least annually and re-anchored to the latest regulatory and threat-landscape evidence. Material changes are version-stamped (see the cover meta block).

Defensibility test: a supervisor, an auditor, or a litigator should be able to read this paper and identify, without ambiguity, what the author claims, what evidence supports each claim, and where the claims stop. That is the institutional standard.

THE CLOSING DOCTRINE

The doctrine in one line.

The institution that relies on penetration testing to evidence its operational resilience is reading the question regulators asked five years ago. Under DORA Article 25, TIBER-EU, and the supervisory crystallisation of these regimes, the question has changed: not "do you have vulnerabilities" but "does your defence wake up under realistic adversarial pressure, and can you evidence it." Adversary emulation answers the second question; nothing else does.

> ***"A pentest finds the holes. An emulation tells you whether the defence woke up. Boards underwrite the second; the regulator credits the second; only the second survives the examination."***

Issued by: Kieran Upadrasta — CISSP · CISM · CRISC · CCSP · MBA · BEng
Affiliations: Schiphol University · Imperials · UCL · ISACA London (Platinum) · (ISC)² London (Gold) · PRMIA · ISF.
Contact: info@kieranupadrasta.com · www.kie.ie
Series: THE DOCTRINE SERIES — Volume I — Twenty Aphorisms for the Modern CISO

CLOSING APHORISM

"A pentest finds the holes. An emulation tells you whether the defence woke up. Boards underwrite the second; the regulator credits the second; only the second survives the examination."

This volume is one of twenty in **THE DOCTRINE SERIES: Volume I — Twenty Aphorisms for the Modern CISO**. Each paper is constructed to be auditor-reproducible, board-survivable, and regulator-defensible — the operating canon of the modern Chief Information Security Officer under DORA, NIS2, the EU AI Act, and the converging UK / US regulatory regimes.

If it cannot be evidenced, it cannot be defended.

Kieran Upadrasta

CISSP · CISM · CRISC · CCSP · MBA · BEng

Cybersecurity Authority · Board Advisor · Interim CISO

www.kie.ie · info@kieranupadrasta.com · linkedin.com/in/kieranupadrasta