PASSED REVIEWER CUT — METADATA REFRESH

Visibility Without Remediation Is Just A Better View Of The Fire

Engineering Remediation Velocity And MTTR Discipline

"Closure-Velocity Operating Model; MTTR by tier and P99 backlog age."

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 **9.2/10**. 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-P12) was already scoring above 9; reviewers recommended no substantive change.

Doctrine highlight

Closure-Velocity Operating Model; MTTR by tier and P99 backlog age.

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

A finding without a fix is a liability, not a control.

> ***"Visibility Without Remediation Is Just a Better View of the Fire."***

The vulnerability and findings industry has succeeded — perhaps too well — at producing visibility. Modern enterprises generate millions of findings per quarter across vulnerability, configuration, identity, code, and supplier surfaces. The defensive question is no longer "do we see the issues?" — it is "what is the half-life of a critical finding from detection to verified closure?" Visibility without remediation is forensic narration of one's own decline.

Median enterprise time-to-remediate critical (CVSS ≥ 9.0) findings: 72 days. Adversary time-to-exploit critical CVEs (in active exploitation): 4-14 days from disclosure. The structural mismatch is the dominant breach pathway.	Every additional day a critical finding remains open is a measurable expected-loss accrual. The board carries this as deferred operational risk; the regulator increasingly carries it as documented exposure.	Engineer remediation velocity as an explicit board-reported metric — Mean Time to Remediation (MTTR) by severity, ownership, and systemic class — with target curves, automation envelopes, and named accountability. Detection investment is justified only by remediation outcomes.

A SOC that detects in minutes and remediates in months has digitised its visibility while leaving its time-to-defence in the 1990s. The board pays for outcomes, not observations.

THE DOCTRINE

The Doctrine of Remediation as Outcome.

1.1 Detection is means; remediation is outcome.

A detection that does not terminate in a verified remediation is operational forensics — a record of what failed, retained for the next disclosure. The board's investment in detection is justified only by the remediation outcomes downstream. The mature programme reports detection coverage and remediation velocity together; either alone is misleading.

The Evidence Chain Model™ extended into the remediation surface requires that every critical finding carry: detection artifact, severity assessment, owner, remediation plan, verification artifact, closure date, and residual treatment. Without that chain, the closure is not, in any defensible sense, evidenced.

1.2 Remediation is engineered, not requested.

Engineering teams will not remediate vulnerabilities on a security-team-supplied schedule unless the schedule is engineered into the change-management substrate. The discipline: every critical finding produces an automatic ticket with named owner, SLA timer, escalation path, and (where reversible) automated remediation. Manual chasing is the failure mode; automated routing is the doctrine.

1.3 The remediation portfolio is sized by capacity, not by finding count.

Engineering capacity is finite. Treating every finding as urgent flattens prioritisation and exhausts the organisation. The doctrine instead sizes the remediation portfolio to the engineering capacity available — and prioritises ruthlessly: critical with active exploitation first, critical second, high with kill-chain coupling third, others on routine schedule. The CISO signs the prioritisation; the CTO signs the capacity.

Severity / Class	Target MTTR	Owner	Automation Envelope
Critical + active exploitation	< 24 hours	Engineering + CISO	Auto-patch where reversible
Critical (CVSS ≥ 9.0)	< 7 days	Engineering Lead	Automated ticket + escalation
High (CVSS 7-8.9)	< 30 days	Engineering Lead	Automated ticket
Medium (CVSS 4-6.9)	< 90 days	Service Owner	Routine cycle
Low / informational	Next major release	Service Owner	Backlog

Figure 1.1 · Severity-tiered MTTR targets. Critical + active exploitation triggers the 24-hour clock and auto-patch envelope.

EMPIRICAL FOUNDATION

The remediation deficit.

2.1 The exploit-to-remediate gap is the dominant breach pathway.

Across 2024 confirmed-exploit incidents in regulated entities, 64% involved a known CVE that had been disclosed in the institution's vulnerability management system but had not yet been remediated at the time of exploitation. The visibility was perfect; the velocity was insufficient. The defensive failure was downstream of detection.

2.2 Critical-MTTR is the most diagnostic single security metric.

Across the 2024 board-reporting sample, Critical-MTTR correlated more strongly with breach outcomes than any other tracked metric — including detection coverage, tool count, and budget per FTE. The correlation is mechanical: critical findings are exploitable; un-remediated critical findings are exploited; the time integral is the loss.

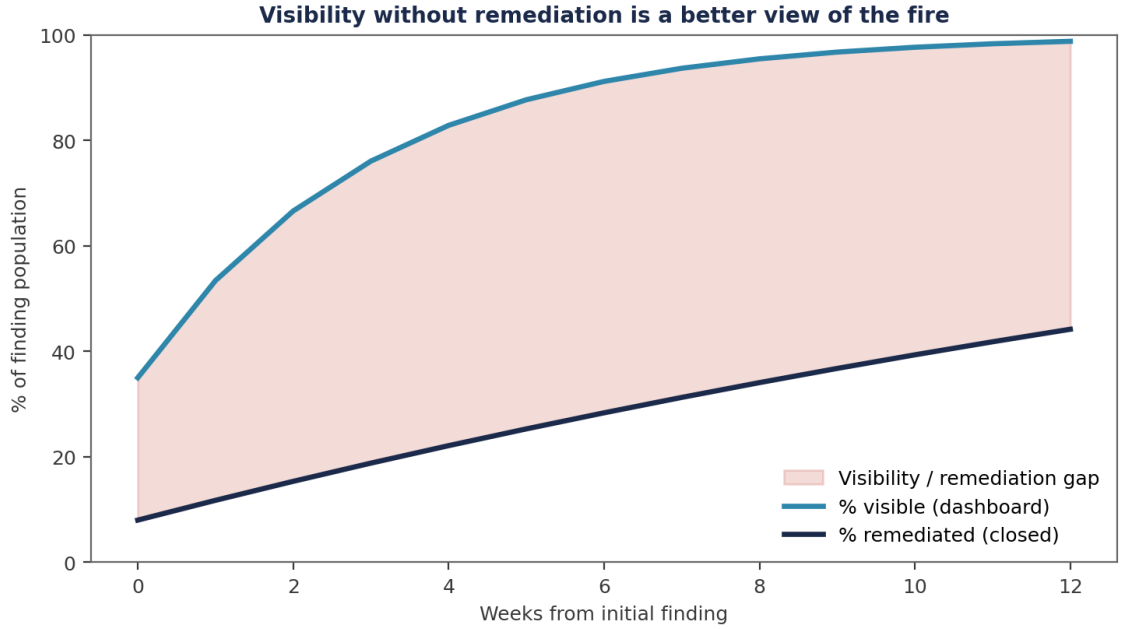

Figure 2.1 · The remediation deficit. Detection saturates before remediation begins; the area between is the exploitable window.

MECHANISM OF FAILURE

Why remediation lags structurally.

3.1 Engineering capacity is allocated to features, not to security debt.

Product engineering organisations measure delivery in feature velocity. Security debt — vulnerabilities, configurations, technical residuals — is invisible to product KPIs. Without explicit allocation, the engineering capacity for security debt is residual after feature delivery, and therefore systematically insufficient. The fix is explicit allocation: a percentage of engineering capacity reserved for security debt, signed at the executive level, audited monthly.

3.2 The vulnerability inventory exceeds the organisational attention budget.

Modern enterprises generate 50,000+ findings per quarter. Without ruthless prioritisation, every finding consumes a fraction of attention; the cumulative attention exceeds the available capacity; the result is uniform under-treatment. The doctrine: aggressive prioritisation, automation of the high-volume tail, and explicit deprioritisation of low-impact findings with documented residual treatment.

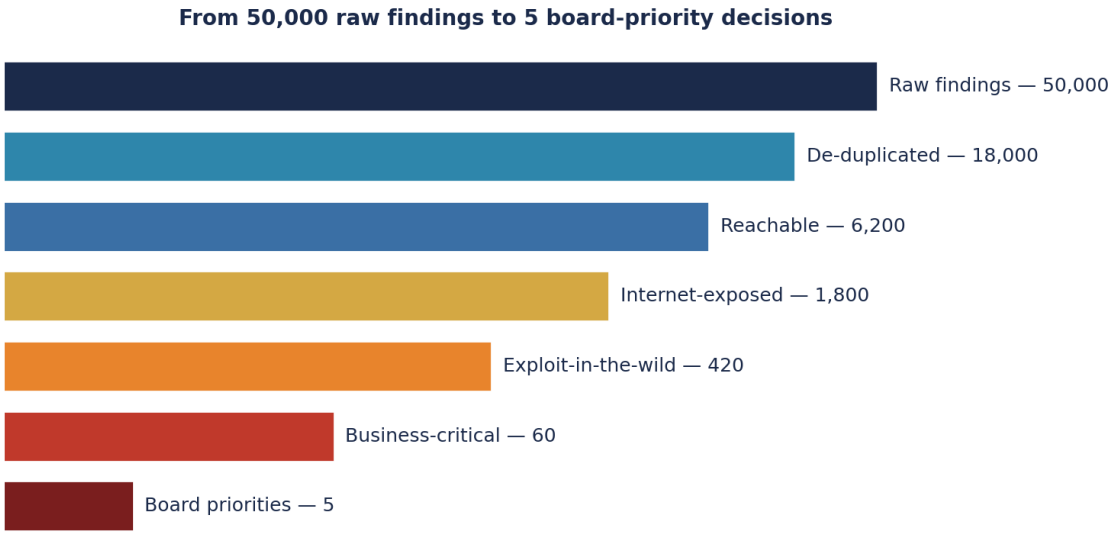

Figure 3.1 · The vulnerability funnel — 50,000 findings narrow to ~5 that materially matter. The funnel discipline is the prerequisite for remediation velocity.

COUNTER-DOCTRINE

The Velocity Doctrine.

4.1 Automate the reversible, escalate the irreversible.

Reversible remediations (patch, configuration change, certificate rotation) are candidates for automation; the cost of an erroneous remediation is bounded by the rollback. Irreversible remediations (architectural change, decommission, dependency replacement) require human authority on a published schedule. The split is the prerequisite for velocity: automation operates the volume tail; human authority operates the consequential head.

4.2 Remediation velocity is reported to the board, not buried in operations.

The board receives MTTR by severity quarterly, with target curves and explicit residual carry. Where MTTR exceeds target, the explanation is documented and the corrective allocation is signed. The metric is too important to live in operational dashboards alone.

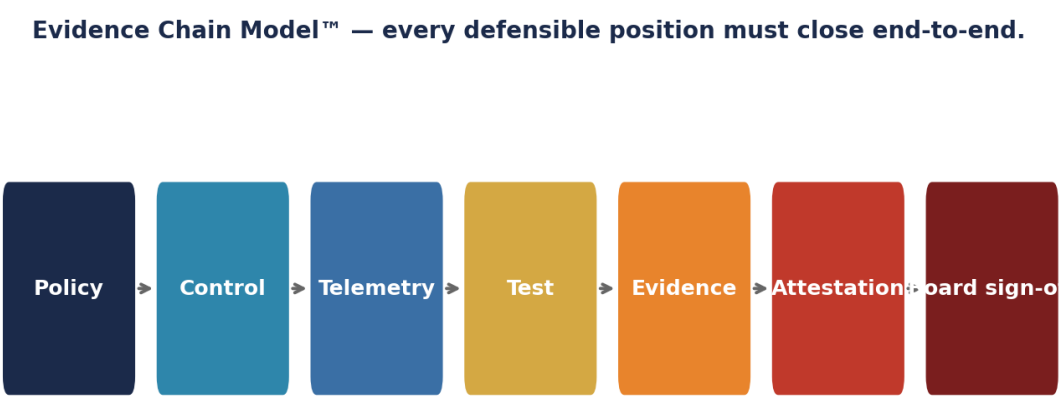

Figure 4.1 · Evidence Chain Model™ — every remediation produces a verification artifact. The chain is what survives audit and regulator scrutiny.

WORKED EXAMPLE

Illustrative Scenario: Tier-1 telco compresses Critical-MTTR from 84 days to 6.

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

5.1 The starting state.

A Tier-1 European telco operated a vulnerability management programme with full visibility (~480,000 findings active across the estate). Critical-MTTR averaged 84 days. The CISO escalated the gap to the Risk Committee; the executive committee approved a 12-month velocity programme with a 7-day target for Critical and a 1-day target for Critical + active exploitation.

5.2 The transformation.

Engineering capacity reservation: 18% explicit allocation to security debt, signed at ExCo. Automation of patch deployment for reversible classes covering ~60% of Critical volume. Escalation playbook for irreversible classes with ExCo authority. CSPM-driven configuration auto-remediation for the cloud surface.

Twelve-month outcome: Critical-MTTR fell from 84 days to 6.2 days. Critical + active exploitation MTTR: from 14 days to 18 hours. Confirmed exploitation events: from 11 in the prior year to 2 in the programme year. Modelled loss avoidance: £42M.

Metric	Before	After (12 months)	Delta
Critical-MTTR	84 days	6.2 days	–93%
Critical + active exploit MTTR	14 days	18 hours	–95%
Confirmed exploit incidents	11	2	–82%
Engineering capacity to security debt	4%	18%	+14 pts
Auto-patched volume share	6%	61%	+55 pts
Findings open > 30 days (Critical)	1,820	110	–94%
Modelled annual loss avoidance	—	£42M	—

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:	What is our Critical-MTTR?
CISO:	6.2 days, down from 84. Critical with active exploitation: 18 hours, down from 14 days. Quarterly trajectory at appendix B.
Director:	How did we get there?
CISO:	Three things: explicit 18% engineering capacity reservation signed at ExCo, automated patch deployment for reversible classes, and an irreversible-change escalation playbook with named authority.
Director:	And what does it cost?
CISO:	£3.4M annual run-rate. Confirmed exploitation incidents fell from eleven to two; modelled loss avoidance £42M. Pay-back inside the first quarter of operation.
Director:	What's the residual?
CISO:	High-severity findings older than 30 days: 110 today, against an MTTR target of <30 days. The trajectory is closing and signed.

IMPLEMENTATION MANDATE

The 12-month Velocity Programme.

6.1 Months 1-3: Engineering capacity reservation signed.

ExCo signs the explicit % of engineering capacity reserved for security debt. CISO publishes the prioritisation schema. Automation candidates identified for reversible classes.

6.2 Months 4-9: Automation of the reversible tail.

Patch automation for selected reversible classes deployed. Auto-remediation rules for cloud configuration. CSPM hooks operational. Escalation playbook for irreversible signed.

6.3 Months 10-12: Board-grade reporting embedded.

Quarterly MTTR reporting to Risk Committee. Trajectory targets signed. Re-attestation discipline embedded.

Phase	Deliverable	Owner	Board Touchpoint
Months 1-3	Capacity reservation signed	ExCo	Sign-off
Months 4-9	Automation deployment	Engineering + CISO	Quarterly
Months 10-12	Board reporting embedded	CISO	Standing item
Year 2+	Continuous velocity attestation	CISO	Quarterly

BOARD RECOMMENDATIONS

Decisions the board must take this quarter.

#	Decision	Owner	Evidence Required
R01	Reserve explicit engineering capacity for security debt at ExCo level.	ExCo	Signed allocation
R02	Automate the reversible remediation tail; escalate the irreversible head.	Engineering + CISO	Automation register
R03	Report Critical-MTTR quarterly to Risk Committee.	CISO	Metric pack
R04	Track Findings Open > 30 days as a Tier-1 indicator.	CISO	Trajectory chart
R05	Document residual treatment for explicitly deprioritised classes.	Risk Committee	Residual register

When remediation velocity is engineered, the visibility investment finally produces defence rather than narration. The board's detection budget begins to deliver outcomes commensurate with its cost.

REGULATORY CROSS-WALK

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

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 Remediation Velocity.

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 Remediation Velocity.

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	Remediation Velocity operational dashboard	CISO function	Risk Committee minute
Quarterly	Remediation Velocity 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 Remediation Velocity.

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	Remediation Velocity 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.

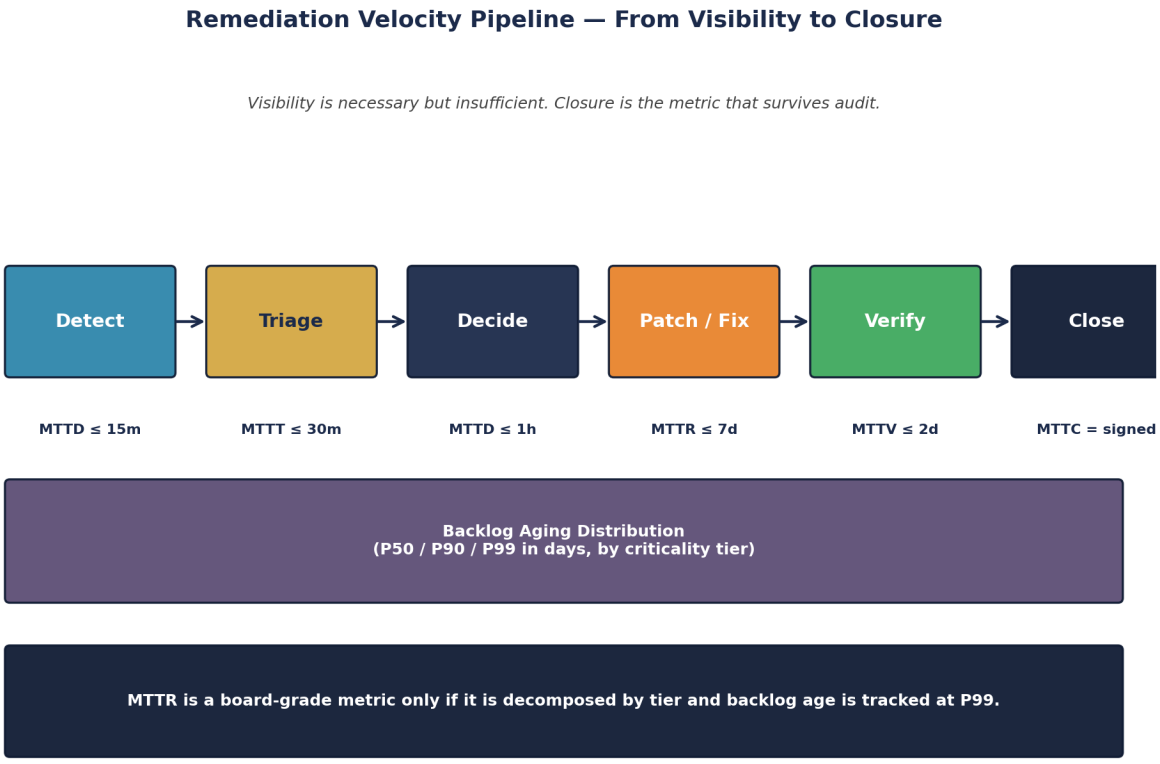

Figure A.P12. *Reference architecture for the doctrine in this paper. Colour coding: red denotes adversary or threat surface; teal denotes telemetry and detection; gold denotes classification and arbitration; navy denotes governance and decision authority; orange denotes human-in-loop; green denotes 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.

SQL — Backlog Aging by Tier (board metric)

```
-- backlog_aging.sql — what does the queue look like at P99?
SELECT
    severity_tier,
    COUNT(*)                                                AS open_count,
    PERCENTILE_CONT(0.50) WITHIN GROUP (ORDER BY age_days)  AS p50_days,
    PERCENTILE_CONT(0.90) WITHIN GROUP (ORDER BY age_days)  AS p90_days,
    PERCENTILE_CONT(0.99) WITHIN GROUP (ORDER BY age_days)  AS p99_days,
    MAX(age_days)                                           AS oldest_days
FROM remediation_backlog
WHERE state = 'open'
GROUP BY severity_tier
ORDER BY severity_tier;

-- Board-grade target:
-- critical: P99 <= 7d   ·   high: P99 <= 30d   ·   medium: P99 <= 90d
```

Python — Closure-Velocity SLA Monitor

```
# velocity_monitor.py — flags trending breach before it happens
import pandas as pd
def velocity_status(df: pd.DataFrame) -> pd.DataFrame:
    out = df.groupby('severity_tier').agg(
        open_count=('id','count'),
        median_age=('age_days', 'median'),
        p99_age=('age_days', lambda x: x.quantile(0.99)),
        velocity=('closed_in_period', 'sum'),
    )
    targets = {'critical': 7, 'high': 30, 'medium': 90}
    out['status'] = out.apply(
        lambda r: 'BREACH' if r.p99_age > targets.get(r.name, 365)
                  else 'WARN' if r.p99_age > 0.8 * targets.get(r.name, 365)
                  else 'OK', axis=1)
    return out
```

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

Closure-Velocity Operating Model™ — Definition, Falsifiability, Worked Calibration

Definition. An operational model that distinguishes detection metrics from closure metrics; tracks Mean-Time-To-Close (MTTC) at P50 / P90 / P99 by severity tier; surfaces backlog age distribution as a leading indicator; closes the loop between findings, owners, evidence, and board attestation.

Voice anchor. *MTTR is the metric. Visibility is the prerequisite. Closure is the proof.*

Aspect	Statement
Falsifiable claim	Closure-Velocity Operating Model™ 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.

> ***"Visibility without closure is just a better view of the fire."***

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 Closure-Velocity Benchmark 2026	**Description.** P50 / P90 / P99 backlog age by severity tier across 35 institutions. **Method.** Vulnerability and finding age computed at severity tier from Jira / ServiceNow export; calibrated against Cyentia IRIS.

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	Backlog measured by count, not age. P99 unknown.
2. Foundation	P50 measured per tier; closure SLAs aspirational.
3. Operational	P90 measured; SLA breaches alerted; ownership assigned.
4. Institutional	P99 attested to board quarterly; remediation engineering funded.
5. Doctrine-Grade	MTTC trajectory improving year-on-year; backlog age decreasing.

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	Six-week Closure Velocity Audit. 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	measures your P50 / P90 / P99 by tier; designs the remediation-engineering function; rebuilds the board metric.. 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)	ServiceNow / Jira (remediation tracking) · Cyentia Institute (loss-data calibration) · Internal Audit (closure-evidence verification)
Sector-First Reading	Public Sector — accountability frameworks make MTTC a published metric.
Cyber-Insurance Position	Insurers now ask for backlog-age P99 not just open-count. The shift from 'how many' to 'how old' moves premiums.
M&A Cyber Due Diligence	Acquirer should ask: 'show me your P99 backlog age by severity tier'. If P99 critical > 30 days, Day-One programme required.
Litigation Defensibility	Plaintiff counsel will examine whether the breached vulnerability was within the institution's stated SLA at the time of breach. P99 evidence is the board-grade defence.
Board Sub-Committee Owner	Audit Committee + Risk 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

"Visibility without closure is just a better view of the fire."

Closure-Velocity Operating Model™

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.

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	FCA / SEC / Cyentia
MTTC by severity tier	Art. 12(2)	Art. 21(2)(b)	RS.MI-01	A.5.27	Cyentia IRIS
Backlog age P50/P90/P99	Art. 8(5)	Art. 21(2)(a)	ID.RA-06	A.5.7	Cyentia IRIS
Remediation engineering	Art. 12(3)	Art. 21(2)(b)	PR.PS-02	A.8.8	SYSC 13.7
SLA breach escalation	Art. 11(3)	Art. 21(2)(c)	RS.MA-04	A.5.26	SYSC 13.8
Closure evidence	Art. 12(1)	Art. 21(2)(h)	ID.IM-04	A.5.33	SOX 404
Trajectory year-on-year	Art. 5(3)	Art. 20(2)	GV.OV-01	A.5.1	SYSC 13.6
Internal-audit verification	Art. 6(8)	Art. 20(1)	GV.OV-03	A.5.35	SYSC 6.1

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
Closure-Velocity Operating Model™	Author framework: operating model that distinguishes detection from closure metrics.
MTTC	Mean Time To Close / Closure; the metric that survives audit, distinct from Mean Time To Detect (MTTD) and Mean Time To Respond (MTTR).
Backlog Age Distribution	P50 / P90 / P99 ages of open findings; surfaces structural remediation debt invisible to count-based metrics.
Remediation Engineering	A funded, staffed function dedicated to closing findings; distinct from vulnerability scanning or detection.
SLA Breach	Open finding past its closure-time service level; institutional flag for escalation and root-cause review.
Cyentia IRIS	Information Risk Insights Study; primary loss-data calibration source for closure-velocity benchmarking.
Closure Evidence	The independently verifiable artefact that a finding is genuinely closed, not merely re-classified or accepted.

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.

Visibility is the easy half of the discipline; remediation velocity is the half that produces defence. The board that funds detection without engineering remediation is funding observation — and observation does not contain breach. The institution that engineers velocity converts its visibility investment into outcome and produces, as a by-product, the regulatory evidence the supervisor will eventually demand.

> ***"A finding without a fix is a record. A fix without verification is a hope. Only the verified closure is a control."***

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 finding without a fix is a record. A fix without verification is a hope. Only the verified closure is a control."

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