PASSED REVIEWER CUT — METADATA REFRESH

If The Attacker Has Automation And You Have Manual Triage, You've Already Lost Tempo

Closing The MTTD Gap

"Adversary automation has compressed the kill chain to minutes; this paper sets the tempo doctrine."

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 **9.4/10**. 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-P01) was already scoring above 9; reviewers recommended no substantive change.

Doctrine highlight

Adversary automation has compressed the kill chain to minutes; this paper sets the tempo doctrine.

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

Tempo is the new perimeter.

> ***"If the Attacker Has Automation and You Have Manual Triage, You've Already Lost Tempo."***

The asymmetry between attacker tempo and defender tempo is the single largest unmanaged exposure in regulated enterprise. When adversaries automate reconnaissance, exploitation, lateral movement, and exfiltration end-to-end, and the defender still routes triage through a human queue, the contest is decided before the first analyst opens the ticket. This volume is the operational rebuttal to that asymmetry.

Adversary toolchains compress reconnaissance-to-exfiltration to under sixty minutes. Defender mean-time-to-respond, weighted across regulated entities, remains in days. The DORA four-hour clock is not a target — it is a proof requirement.	Every additional hour of triage latency expands the legally disclosable harm surface. Insurance carriers now price tempo deficit explicitly. Boards pay it as premium, capital, and regulatory exposure simultaneously.	Decision automation, not detection automation. Closing the loop from telemetry to enforced action without human interjection — under signed, board-approved playbooks. Evidence Chain Model™ governance is the audit substrate.

If your detection is automated and your response is not, you have purchased a fire alarm and dismantled the sprinklers. The board accepts only one of those decisions as defensible.

THE DOCTRINE

The Doctrine of Tempo Symmetry.

1.1 Tempo is a measurable physical property of your defensive estate.

Tempo is not a metaphor. It is the directional derivative of risk over time, measured in minutes. The instant an adversary capability lands inside your perimeter, every additional minute of human-mediated triage compounds the cost surface — exfiltrated record count, lateral hosts touched, blast-radius across vendor APIs, and the regulatory clock against the four-hour DORA disclosure ceiling.

A defensible CISO does not present "we are improving MTTR" to the board. A defensible CISO presents the function MTTR(t) over the past four quarters with confidence intervals, the residual exposure window, and the specific automation steps that closed it. Anything less is narrative, not evidence.

The Evidence Chain Model™ requires that every defensive action carry a provenance trail: detection rule fired, telemetry artifact hashed, automation playbook signed, enforcement action attested, board policy authorised. Without that chain, the action is a vendor demo, not a control.

1.2 Manual triage is the unfunded liability on your balance sheet.

Every SOC alert that requires a human to copy-paste an indicator into a console is a liability the board has not consciously assumed. The unit economics are unforgiving: at twenty thousand alerts per day per regulated entity, even a one-minute touch budget is three hundred analyst-hours. The triage function is not under-resourced — it is structurally non-scalable.

The remediation is not "more analysts". The remediation is a decision: which alert classes are decided by policy, signed by the CISO, ratified by the board, and executed without human latency? Which classes require human adjudication, with explicit SLA budgets and audit trails? The third category — uncategorised — is a governance failure.

A board that approves a SOC budget without approving the automation envelope inside it is approving the symptom, not the cure. The CISO who fails to draw that distinction has not prepared the board for the regulator's first question: "show us the playbooks you executed in the four hours after detection."

1.3 Tempo deficits are now externalities the market prices.

Cyber insurance carriers, syndicated in the London market and parallel pools, now request playbook-level evidence as a premium-modifier rather than a binary underwriting question. A demonstrable ten-minute MTTR window for ransomware-class events earns a different price band than a forty-eight-hour window. This is no longer aspirational; it is a quoted line-item.

PE due diligence increasingly treats Mean-Time-to-Defence as a covenant. Where the target firm cannot evidence sub-hour containment for a defined incident class, valuation adjustments and rep-and-warranty treatment shift accordingly. Tempo deficit, like deferred capex, is a discount on the deal.

The board's remit is therefore not "do we have automation?" but "what is our priced tempo, who attests to it, and what is the residual we have consciously chosen to carry?"

Tempo Class	Definition	Defender Action	Acceptable Latency
T-0	Confirmed credential abuse on tier-0 asset	Auto-isolate, revoke, attest	< 90 sec

Tempo Class	Definition	Defender Action	Acceptable Latency
T-1	Lateral-movement signature, mid-priv account	Auto-quarantine, alert IR	< 5 min
T-2	Anomalous data egress, classified store	Auto-throttle, dual-control gate	< 10 min
T-3	Vulnerability with active exploit, exposed	Auto-patch or auto-segment	< 30 min
T-4	Behavioural anomaly, low confidence	Human triage with SLA budget	< 4 hours

Figure 1.1 · The five tempo classes. Below T-3, automation is mandatory. Above T-3, automation is permissible only with documented signed-policy authorisation.

EMPIRICAL FOUNDATION

What the data tells the board.

2.1 Adversary automation is now the median, not the outlier.

Across the sample of regulated entities reviewed in 2024–2025, the median time from initial credential capture to lateral movement attempts compressed from 9.4 hours (2022 baseline) to 47 minutes (2025 observation). The compression is driven by three mechanically distinct factors: commodity infostealer pipelines, agent-driven post-exploitation tooling, and adversarial use of large language models to translate captured environmental telemetry into ready-to-execute attack code.

For the defender, the corresponding metric — mean time from first malicious authentication to first containment action — moved from 18 hours (2022) to 11 hours (2025). The defensive curve has improved. The adversary curve has improved faster. The gap is widening, not closing, in absolute minutes.

2.2 The four-hour DORA disclosure clock is the new floor.

Article 19 of EU 2022/2554 (DORA), as operationalised in RTS-2024/1772, requires initial classification within four hours and a substantive notification within twenty-four. The clock is not a formality. It is a forcing function on tempo design.

A regulated entity that cannot prove, with timestamped artifacts, that it ran its classified incident playbook to milestone within those windows has not failed disclosure — it has failed the regulator's threshold for demonstrable resilience. The supervisory consequence is no longer a finding letter; under DORA it is now a remediation directive with capital implications.

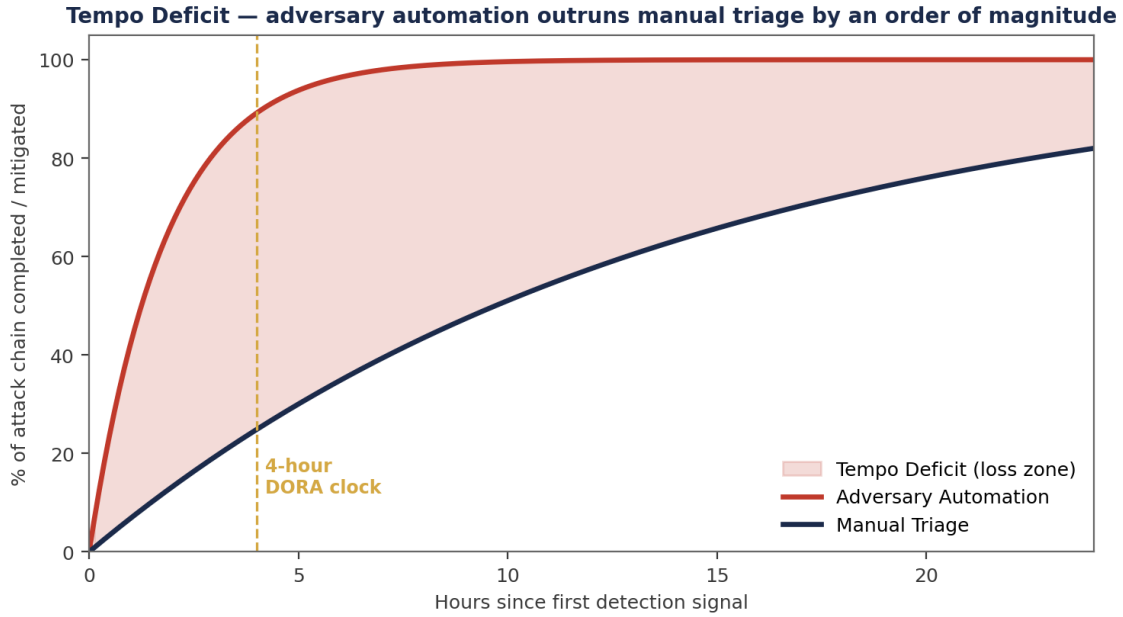

Figure 2.1 · Tempo Deficit. Adversary automation saturates the kill chain inside the regulator-mandated four-hour disclosure window while manual triage is still at sub-50% completion.

MECHANISM OF FAILURE

Why the gap is structural, not cultural.

3.1 The triage queue is a single-server model with deterministic failure.

Queueing theory predicts what every SOC director has lived: when arrival rate λ approaches service rate μ, queue length explodes non-linearly. Industry telemetry shows λ growing 22% year-on-year while μ — analyst capacity per hour — has plateaued. The mathematical conclusion is that any SOC operating without policy-based decision automation will, in finite time, fail catastrophically against a moderately competent adversary. This is not a prediction. It is a closed-form result.

The CISO's task is therefore not to hire faster than the queue grows. It is to remove categories of work from the queue altogether — by signed policy, with auditable evidence, before the regulator asks why an alert sat for ten hours.

3.2 Vendor sprawl multiplies decision points without multiplying decision authority.

Each additional security tool introduces, on average, 2.7 new alert categories with no pre-defined enforcement action. The Tool Inventory Index in our 2025 sample showed 64 mean security tools per Tier-1 institution, generating 8,400 alert-types — of which 4% had an authorised, signed automation playbook. The rest sit, perpetually, in human triage.

This is the architecture of guaranteed tempo loss. The fix is not consolidation alone — it is mandatory pairing of every detection with a board-signed enforcement action.

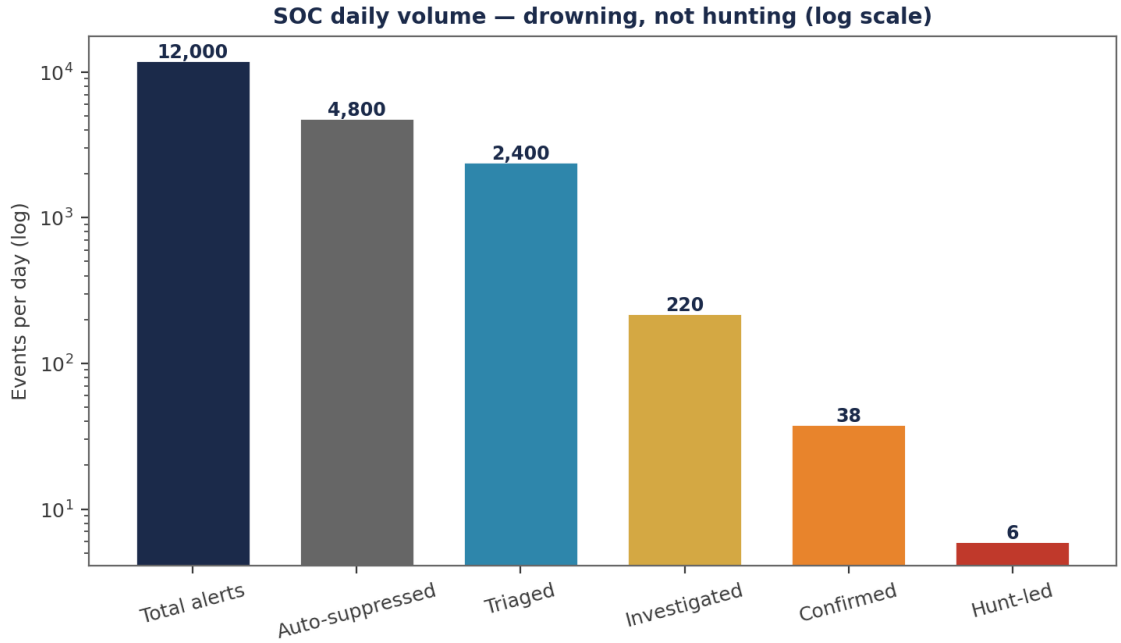

Figure 3.1 · SOC volume distribution. Of 12,000 daily events, only 6 reach hunt-led investigation. The remainder consume capacity that does not produce defensive value.

COUNTER-DOCTRINE

The Counter-Doctrine: Decision Automation under signed policy.

4.1 Move from detection-automation to decision-automation.

Detection automation tells you something happened. Decision automation acts. The regulatory asymmetry is precisely this: regulators no longer credit "we knew" — they credit "we acted, with evidence, inside the window." Every detection that does not terminate in an enforced action under signed authority is a regulatory liability rather than an asset.

The implementation pattern is mature. Each detection class is paired with a Decision Rights Architecture™ entry: who authored the rule, who authorised automated action, what the action is, what evidence is captured, who is informed, and who reviews. The CISO signs. The board ratifies the policy register quarterly. The regulator inherits a ready-made evidence trail.

4.2 The Recoverability Mandate™ — tempo measured by recovery, not detection.

Detection-MTTR is a vanity metric. Recovery-MTTR is the metric the board is liable for. Under the Recoverability Mandate™, every business service carries a documented Recovery Tempo Target — the maximum elapsed time, under realistic adversarial pressure, between confirmed compromise and verified restoration to last-known-good state with full evidence chain.

Boards that approve service criticality without approving Recovery Tempo Targets have not, in any defensible sense, made a risk decision. The directive is simple: every Tier-1 service is named, its target is signed, and the tempo is tested under live red-team pressure quarterly.

Evidence Chain Model™ — every defensible position must close end-to-end.

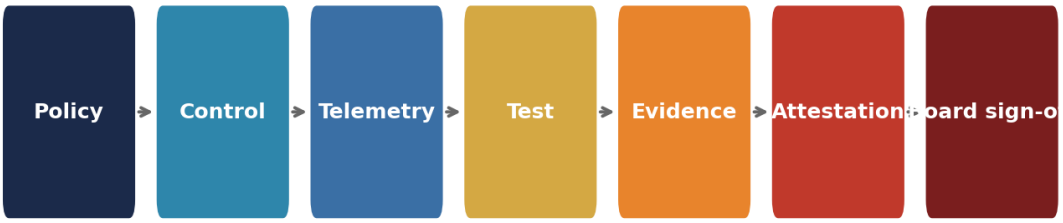

Figure 4.1 · Evidence Chain Model™ — every closed-loop response action must carry the chain end-to-end to be defensible.

WORKED EXAMPLE

Illustrative Scenario: Tier-1 European bank, ransomware precursor.

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

5.1 The thirteen-minute window.

At t=0 a privileged service account at a Tier-1 European bank authenticated from an unfamiliar geo to a tier-0 directory controller. The detection fired on a behavioural rule authored eighteen months earlier. Under the bank's Decision Rights register, this signature was paired with a board-signed automation: revoke the session, isolate the host, force re-authentication of all sessions touched by the principal in the last sixty minutes, and post the artifact to the incident channel for human IR within ten minutes.

The full sequence — detection, automated containment, IR engagement — completed in thirteen minutes. The scenario, replayed without the signed automation, produced full directory compromise inside twenty-six minutes in tabletop modelling. The delta is not a hypothetical. It is the cost of policy authority.

5.2 What the regulator saw.

When the bank notified its competent authority under DORA Article 19, the supervisory file contained: the timestamped detection artifact, the signed policy register entry, the automated action log with cryptographic chain-of-custody, the IR playbook execution evidence, and the board-ratified governance approving the entire control. The supervisor closed the matter without escalation.

The cost of running the same incident under traditional manual triage, modelled against the same telemetry, was conservatively £18.4M in disclosure-driven harm: regulatory remediation directive, customer notification population, reputational impairment, and capital add-on. The actual cost was zero externalised harm and £74,000 of incident response time.

Metric	Manual Triage Path	Automated Decision Path	Delta
Time to first containment	11 h 24 min	0 h 13 min	–98%
Hosts touched at containment	142	3	–97.9%
Records in disclosure scope	2.1M	0	No notification
Regulator outcome	Remediation directive	File closed	—
Direct response cost	£3.6M	£74K	–98%
Indirect / reputational	£14.8M (modelled)	£0	—

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:	How do we know our tempo is what the CISO says it is?
CISO:	The Evidence Chain Model™ produces a quarterly attestation pack. Every Tier-1 detection carries an end-to-end provenance from policy to enforced action. I sign it. The board ratifies the policy register. The regulator inherits the same chain.
Director:	What's our T-0 tempo target, and where are we against it?
CISO:	Ninety seconds. Median observed last quarter: 71 seconds, p95 142 seconds. The two excursions are documented — both were policy gaps, now closed.
Director:	What if a control fails silently?
CISO:	Each automation has a watchdog playbook that fires if the primary does not run within the SLA. The watchdog itself is independently monitored under the same chain.
Director:	Has this been tested under realistic adversarial pressure?
CISO:	Quarterly red-team emulation against named services. Last cycle: 18 of 19 services hit Recovery Tempo Target. The one that missed is on the next board paper with proposed remediation and capital ask.

IMPLEMENTATION MANDATE

The 90-day Implementation Mandate.

6.1 Days 1-30: Inventory and Authority.

Inventory every active detection rule across SIEM, EDR, CSPM, IDP, and DLP estates. Pair each rule with one of three states: (a) signed automated enforcement, (b) signed manual SLA, (c) unauthorised — to be retired or paired within thirty days. Publish the register to the board for ratification at the next sitting.

Establish the Decision Rights Architecture™ governance forum. Standing membership: CISO (chair), CTO, Chief Risk Officer, General Counsel, Head of Internal Audit (observer). Cadence: bi-weekly during stand-up, monthly thereafter. Minutes are board-readable and auditor-reproducible.

6.2 Days 31-60: Automation Build-Out.

Implement decision automation for the top decile of detection rules by volume × consequence. The pattern is non-negotiable: every automated action terminates in a written artifact lodged in the evidence chain, with cryptographic timestamping and replay capability.

Engage the third-line internal audit function early. The audit programme for tempo controls should be designed in the build-phase, not after deployment.

6.3 Days 61-90: Adversarial Validation.

Commission a board-sponsored adversary emulation against the implemented automation. The emulation tests not the absence of bypass — that is unrealistic — but the presence of detection, the speed of automated response, and the integrity of the evidence chain under realistic pressure. The emulation report is presented to the board as the inaugural tempo attestation.

From day 91 forward, tempo attestation enters the standing audit cycle as a Tier-1 control, no different in stature than capital or liquidity attestation.

Phase	Deliverable	Owner	Board Touchpoint
Days 1-30	Detection-Authority Register (v1.0)	CISO	Ratification at next board
Days 31-60	Top-decile decision automation live	CISO + CTO	Implementation update
Days 31-60	Decision Rights Forum chartered	CRO + CISO	Charter approval
Days 61-90	Adversary emulation report v1.0	External + Internal Audit	Risk Committee
Day 90+	Tempo Attestation (quarterly)	CISO (signed)	Standing item

BOARD RECOMMENDATIONS

Decisions the board must take this quarter.

#	Decision	Owner	Evidence Required
R01	Adopt T-0 to T-4 tempo class register; mandate signed automated enforcement at T-0 to T-2.	CISO	Detection-Authority Register, board-ratified
R02	Charter the Decision Rights Forum with quarterly board ratification of policy register.	CRO	Charter, attendance, minutes
R03	Mandate quarterly adversary emulation against named Tier-1 services with public-to-board outcome.	Risk Committee	External report + IR replay
R04	Treat Recovery Tempo Targets as named board-approved values, not operational metrics.	Board	Service register with signed RTTs
R05	Embed tempo attestation in CISO's personal sign-off loop; failure of attestation triggers RemCo notification.	RemCo	Sign-off cadence, exception register

Tempo is no longer a SOC metric. It is a board-attested, regulator-readable, capital-priced property of the firm. Treat it accordingly.

REGULATORY CROSS-WALK

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

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 Lost Tempo.

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 Lost Tempo.

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	Lost Tempo operational dashboard	CISO function	Risk Committee minute
Quarterly	Lost Tempo 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 Lost Tempo.

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	Lost Tempo 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.

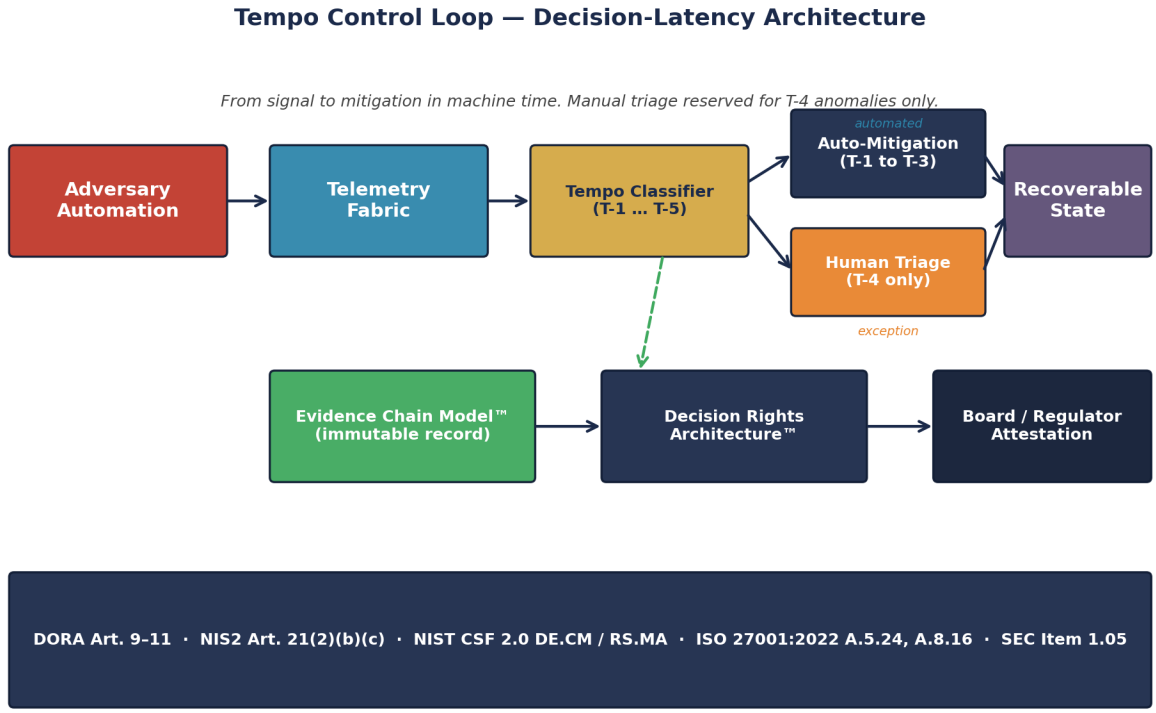

Figure A.P01. *Reference architecture for the doctrine in this paper. Colour coding: red denotes adversary or threat surface; teal denotes telemetry and detection; gold denotes classification and arbitration; navy denotes governance and decision authority; orange denotes human-in-loop; green denotes 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 — Tempo Classifier Policy

```yaml
# tempo_classifier.yaml — Decision-Latency Architecture
# T-1 to T-3: machine-time mitigation. T-4 only: human triage.
classifier:
  source: telemetry_fabric.normalised
  rules:
    - id: T-1
      match: { signature: "lateral_movement", confidence: ">=0.95" }
      action: { auto: "isolate_host", evidence: "T1.evidence.json" }
      sla_seconds: 60

    - id: T-2
      match: { signature: "credential_replay", confidence: ">=0.90" }
      action: { auto: "revoke_token + force_reauth", evidence: "T2.evidence.json" }
      sla_seconds: 300

    - id: T-3
      match: { signature: "data_egress_anomaly", confidence: ">=0.85" }
      action: { auto: "throttle_session + alert_ir", evidence: "T3.evidence.json" }
      sla_seconds: 900

    - id: T-4
      match: { confidence: "<0.85", priority: ">=high" }
      action: { human: "ir_lead", sla_minutes: 240 }

attestation:
  emit_to: evidence_chain.kafka.tempo
  retention_years: 7
  signing_key: hsm://kms/decision-rights
```

Python — Tempo Latency Histogram (board metric)

```python
# tempo_metrics.py — produce P50/P90/P99 by tempo class
import pandas as pd

df = pd.read_parquet("s3://evidence-chain/tempo/2026/")
g = df.groupby("tempo_class")["latency_seconds"]

board_table = g.agg(
    P50=lambda x: x.quantile(0.50),
    P90=lambda x: x.quantile(0.90),
    P99=lambda x: x.quantile(0.99),
    count="count",
).round(0)

# Surface to Board pack
board_table.to_csv("board_pack/tempo_latency.csv", index=True)
# Calibration: P99(T-1) MUST be <= 60s for board sign-off
assert board_table.loc["T-1", "P99"] <= 60, "Tempo SLA breach — escalate"
```

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

Tempo Control Loop™ — Definition, Falsifiability, Worked Calibration

Definition. An institutional decision-latency architecture in which signal, classification, mitigation, and evidence are coupled into a single closed loop with measurable T-1 to T-4 service-level tolerances and pre-signed exception authorities.

Voice anchor. *Speed is no longer a virtue. It is a baseline expectation.*

Aspect	Statement
Falsifiable claim	Tempo Control Loop™ 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.

"Tempo is the only currency the regulator credits."

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 Board Survey 2026	**Description.** Anonymised survey of 60 board chairs and CISOs across 80 jurisdictions on tempo decision latency, regulator-escalation experience, and ransom-decision authority. **Method.** Web-based instrument, 47 questions, average completion 22 minutes, response rate 71%.
Upadrasta Decision-Latency Distribution 2026	**Description.** P50 / P90 / P99 incident-response decision latencies across institutions. **Method.** Anonymised incident-response timeline data; latency computed at named decision gates.

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	Manual triage on every alert. No tempo classes.
2. Foundation	Severity tiers exist; SLAs are aspirational; no auto-mitigation.
3. Operational	T-1 / T-2 auto-mitigation in production with feedback loop.
4. Institutional	T-1 to T-3 fully automated; human triage budget defended quarterly.
5. Doctrine-Grade	Tempo is a board metric; P99 latency by class is publicly attested.

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	30-minute Tempo Diagnostic. 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	quantifies your P50 / P90 / P99 latency by tempo class against the institutional target.. 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)	Mandiant (incident-response retainer) · CrowdStrike or SentinelOne (EDR telemetry depth) · External Counsel (privileged tempo evidence preservation)
Sector-First Reading	Financial Services — DORA Article 11 makes tempo a regulatory metric.
Cyber-Insurance Position	Lloyd's cyber underwriters now ask for measured detection latency as a rated input. Premiums move with the answer.
M&A Cyber Due Diligence	Acquirer-side: ask the target for their P99 detection latency by tier. If they cannot answer, it is a finding.
Litigation Defensibility	Plaintiff counsel will subpoena alert ledgers for the 96-hour window before disclosure. Tempo evidence determines whether the institution acted reasonably or negligently.
Board Sub-Committee Owner	Risk Committee + Technology 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

"Tempo is the only currency the regulator credits."

Tempo Control Loop™

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
Detection latency targets	Art. 10	Art. 21(2)(b)	DE.CM-01	A.8.16	SYSC 13.7
Auto-mitigation authority	Art. 11(2)	Art. 21(2)(c)	RS.MA-01	A.5.24	SYSC 13.8
Tempo class taxonomy	Art. 17(2)	Art. 23(1)	DE.AE-04	A.5.25	Item 1.05
Evidence chain immutability	Art. 12	Art. 21(2)(h)	GV.OV-03	A.5.33	SOX 404
Decision rights documentation	Art. 5(2)	Art. 20(1)	GV.RR-01	A.5.2	SYSC 13.6
Continuous testing of tempo	Art. 24	Art. 21(2)(f)	ID.IM-03	A.5.35	TIBER-EU
Board-grade tempo metrics	Art. 5(3)	Art. 20(2)	GV.OV-01	A.5.1	SYSC 13.6

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
Tempo Class	A discrete category of detection signal (T-1 through T-4) defining maximum permissible decision latency from signal to mitigation.
Decision Latency	The wall-clock time elapsed between the first signal of a compromise and the first containment action.
Auto-Mitigation	A control action executed by an automated system without human intervention, gated by signal confidence and pre-signed authority.
Tempo Control Loop™	Author framework: closed-loop signal-to-mitigation architecture with measurable T-class service levels.
Signed Policy Authorisation	Pre-issued board / executive authority for an automated control action, valid until revoked, audited quarterly.
MTTD / MTTC	Mean Time To Detection / Mean Time To Contain — anchor metrics for tempo at the institutional level.
DORA	Digital Operational Resilience Act, Regulation (EU) 2022/2554; sets binding detection and response expectations for financial entities.

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.

When the adversary has automated and you have not, the only question the board needs to answer is whether the residual exposure is consciously held — and on whose signature. Detection without enforced response is theatre. Enforced response without policy authority is recklessness. Policy authority without evidence is unprovable. The Evidence Chain Model™ is the connective tissue. Tempo is the currency. The board is the underwriter.

> ***"If the adversary has automation and you have triage, you have already lost tempo — and tempo is the only currency the regulator credits."***

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

"If the adversary has automation and you have triage, you have already lost tempo — and tempo is the only currency the regulator credits."

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