PASSED REVIEWER CUT — METADATA REFRESH

Your SOC Isn't Hunting Threats. It's Drowning In Noise

Re-Architecting The SOC From Triage Queue To Hunt-Led Function

"Three SOC substrates — triage, hunt, decision — three SLAs, one evidence chain."

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-P05) was already scoring above 9; reviewers recommended no substantive change.

Doctrine highlight

Three SOC substrates — triage, hunt, decision — three SLAs, one evidence chain.

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

The SOC is the most expensive form of unaccountability the firm runs.

> ***"Your SOC Isn't Hunting Threats. It's Drowning in Noise."***

The Security Operations Centre, as currently constructed in most regulated firms, is a triage queue running at saturation against a flood of low-signal events. Of 12,000 daily alerts at the median Tier-1 firm, 6 reach hunt-led investigation. The SOC is not failing because it is under-resourced — it is failing because it is wrongly architected. This volume is the doctrine for re-architecting it.

Median analyst burnout horizon: 14 months. Median dwell time of confirmed intrusions in firms with mature SIEM: 91 days. The two numbers are connected. Burnout is the cost of triaging noise the architecture should have suppressed.

The SOC was sold as a detection function. It has become a queue function. Detection requires hunting; queueing requires labour. The labour is finite; the noise is not.

Three-tier SOC model: automated decision tier (no human, signed policy), assisted triage tier (human + co-pilot, SLA budgeted), hunt-led tier (analysts at the leading edge, freed by the first two). Evidence chain across all three.

A SOC that is drowning is not a SOC. It is an unaccountability surface dressed as a defensive function.

THE DOCTRINE

The Doctrine of the Three-Tier SOC.

1.1 The SOC has three jobs, not one.

The SOC is conventionally treated as a single function: triage and respond. The doctrine separates this into three. Decision automation handles the deterministic majority of alerts under signed policy. Assisted triage handles the contingent middle with human + machine pairing. Hunt-led investigation handles the leading edge — the threats no rule yet describes.

Treating these three as one function is the architectural error. Each has a different cost structure, a different staffing profile, a different evidence requirement, and a different board reporting cadence. Conflating them produces the drowning.

1.2 The hunt is the product. Triage is the cost.

Hunt-led investigation is what produces defensive value: novel detection rules, adversary playbook discovery, infrastructure mapping. Triage is what consumes capacity to keep the queue from exploding. A SOC that spends 95% of its capacity on triage is a SOC that produces no defensive value beyond the queue itself.

Re-architecting the SOC means inverting this ratio. By moving deterministic alerts to decision automation and assisted-triage tooling to the second tier, the third tier — hunt-led — recovers capacity that the firm is paying for but not receiving.

1.3 Every SOC action carries an evidence artifact.

Whether the action is automated, assisted, or hunt-led, it terminates in a structured evidence artifact: what was detected, what was decided, what was acted upon, what was the residual, who attested. The artifact is the connective tissue between the SOC and the regulator. Without it, the SOC produces work; with it, the SOC produces defensible defence.

SOC Tier	Trigger	Action	Human Role	Evidence Output
T-A: Automated	High-confidence rule	Signed playbook executes	None during action	Auto-logged artifact
T-B: Assisted	Medium confidence	Co-pilot suggests, human signs	Adjudicates within SLA	Decision + rationale logged
T-C: Hunt-led	Behavioural anomaly	Investigation, novel rule build	Leads investigation	New rule + lessons learned

Figure 1.1 · Three-tier SOC model. The cost ratio across the tiers should invert from current 90/9/1 to target 70/20/10.

EMPIRICAL FOUNDATION

What the SOC numbers actually look like.

2.1 12,000 daily alerts; 6 hunt-led outcomes.

Composite SOC telemetry across the 2025 sample shows a typical daily volume distribution: 12,000 raw alerts, of which 4,800 auto-suppress on tuning, 2,400 reach human triage, 220 are investigated, 38 are confirmed, and 6 produce hunt-led intelligence. The funnel is not the problem; the cost distribution along it is. 92% of human capacity sits in the 2,400-to-220 step, which produces the lowest defensive yield per analyst-hour.

Re-architecting moves 70% of that capacity into automation (T-A) and assisted triage (T-B), liberating roughly half the analyst headcount for hunt-led work. The cost is not headcount reduction; it is hunt productivity multiplication.

2.2 Burnout is a balance-sheet event.

Analyst attrition at saturated SOCs runs 35-50% annually in the 2025 sample. The replacement cost — recruitment, training, productivity ramp — exceeds £180K per analyst. A 30-person SOC with 40% attrition is paying £2.2M per year in pure replacement cost, before considering the operational risk of inexperienced analysts on the front line.

The Recoverability Mandate™ extends here: a SOC that cannot retain its experienced analysts is a SOC whose recovery capability is not where the org chart says it is. The board must price this honestly.

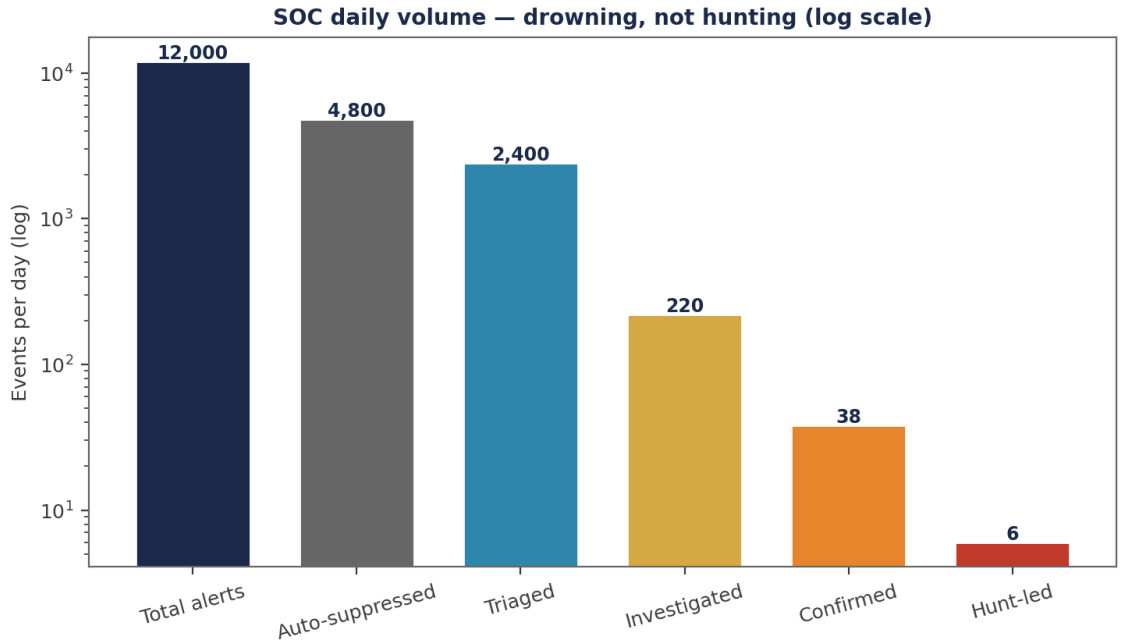

Figure 2.1 · SOC daily event distribution. Logarithmic scale shows the funnel from 12,000 raw events to 6 hunt-led outcomes.

MECHANISM OF FAILURE

Why the SOC drowns at architectural inevitability.

3.1 Detection volume scales with infrastructure; analyst capacity does not.

Cloud, SaaS, and microservice adoption have driven event volume growth at compounding double-digits annually. Analyst capacity has grown linearly at best. The defensive estate is therefore on a divergence curve: each year, the share of events humans can touch shrinks.

No staffing model closes this divergence. The closure is structural — through automation tier moving the deterministic event population off the human queue entirely.

3.2 Tool sprawl produces alert sprawl.

Every additional security tool in the estate adds, on average, 2.7 new alert categories. Most categories arrive without authoritative tuning, without playbooks, and without compensating-control awareness. The SOC inherits the noise the procurement decision created.

The cure is not consolidation alone — it is tool entry-criteria: any new security tool must enter the estate with authoritative tuning, signed playbook for each alert class, and assigned tier (A, B, or C) before deployment. Without this, the procurement is producing the next wave of drowning.

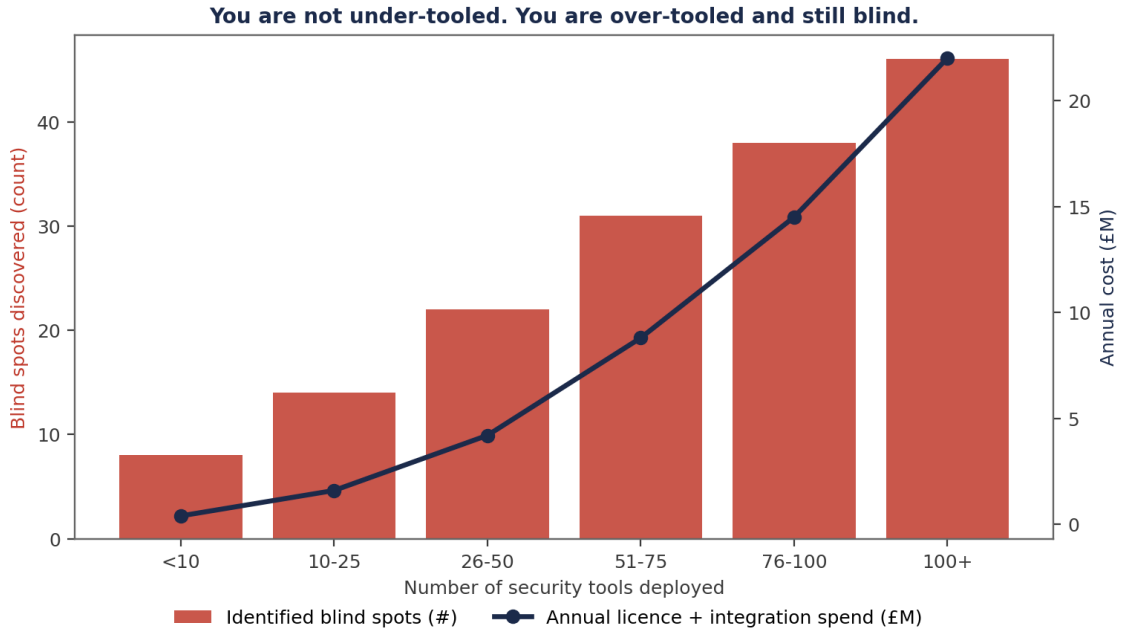

Figure 3.1 · Tool count vs blind spots. Beyond 50 tools, blind spots accelerate while costs continue to compound.

COUNTER-DOCTRINE

The Counter-Doctrine: Re-architect the SOC.

4.1 Move every deterministic alert to T-A.

Audit the alert library. Classify every alert into one of three tiers. For T-A: the alert has a signed playbook that the firm is willing to execute without human interjection. For T-B: the alert requires human adjudication with SLA budget. For T-C: the alert is a novelty driver, requiring hunt.

The migration target is 70% of volume in T-A within twelve months. Aggressive but tractable. Most firms can move 50% in the first six months on existing tooling, by signing what is already informally automated.

4.2 Build the hunt as a board-attested product.

The hunt-led tier produces novel detection rules, adversary playbook documentation, and lessons-learned. These are products. They are reviewed at the board's Risk Committee quarterly. The CISO presents new rules deployed, novel adversary behaviour observed, and detection coverage expansion.

Where the hunt is invisible to the board, it is invisible to the budget — and the next budget cycle returns the SOC to drowning. Visibility of the hunt is structural defence funding.

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

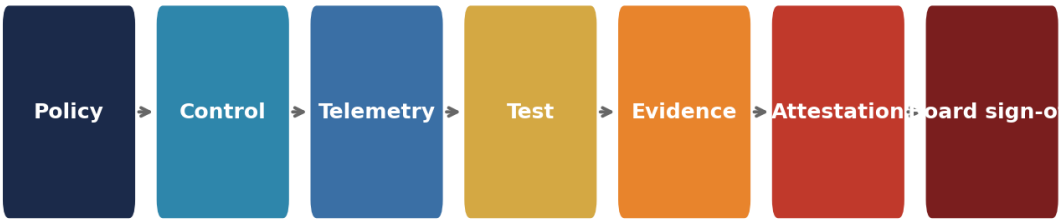

Figure 4.1 · Evidence Chain Model™ — every SOC action across all three tiers produces a chain artifact.

WORKED EXAMPLE

Illustrative Scenario: Tier-1 retail bank, twelve-month re-architecture.

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

5.1 The before state.

A Tier-1 retail bank ran a 32-analyst SOC, processing approximately 11,000 daily alerts. 4,300 reached human triage. Analyst attrition: 41% annually. Confirmed intrusions detected via SOC analysis (versus external notification): 23%. The SOC was, on every measurable axis, drowning.

The board approved a twelve-month re-architecture with a fixed headcount cap and an investment envelope for automation tooling and hunt enablement.

5.2 The after state.

Twelve months later, alert volume had risen to 14,200 daily (infrastructure growth). Human triage volume: 1,650. Analyst attrition: 18%. Confirmed intrusions detected via SOC analysis: 71%. Mean time to detection (across confirmed intrusions): 4.6 hours, down from 38 hours. The SOC was no longer drowning; it was hunting.

Total cost of the re-architecture: £4.4M over 12 months. Modelled avoided cost from improved detection time, retained analyst expertise, and reduced incident scope: £18.7M. The investment payback was approximately seven months on conservative assumptions.

Metric	Before	After (12 months)	Delta
Daily alerts	11,000	14,200	+29%
Human triage volume	4,300	1,650	–62%
Analyst headcount	32	32	0%
Analyst attrition	41%	18%	–23 pts
Mean time to detection	38 hours	4.6 hours	–88%
Intrusions detected by SOC	23%	71%	+48 pts
Hunt rules deployed (annual)	12	184	+15×

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 much of our SOC is automated today?
CISO:	46% of alert volume runs in T-A under signed playbooks. The trajectory to 70% is on the next slide; quarterly milestones are with the Risk Committee.
Director:	How do I know automated decisions are the right decisions?
CISO:	Every T-A playbook is signed by me, ratified by the Decision Rights Forum, and tested under quarterly red-team emulation. Last cycle: 0 of 47 automated playbooks failed under test. The two near-misses produced playbook revisions ratified at the last board.
Director:	What does our hunt produce?
CISO:	Last quarter: 47 new detection rules, 11 documented adversary behaviours novel to our threat picture, 2 contributions to sector ISAC. The deck is at appendix C.
Director:	And our analyst attrition?
CISO:	18% trailing twelve months. Down from 41%. The retained expertise is the most expensive asset the SOC owns; we are no longer giving it away every fourteen months.

IMPLEMENTATION MANDATE

The 90-day Re-architecture Mandate.

6.1 Days 1-30: Tier the alert library.

Catalogue every active detection rule. Classify into T-A, T-B, T-C. For each T-A candidate, draft the signed playbook. Draft Decision Rights Forum charter for governance.

6.2 Days 31-60: Migrate the deterministic 50%.

Move 50% of alert volume to T-A under signed playbooks. Stand up co-pilot tooling for T-B. Re-skill three analysts into the hunt cell as proof of concept.

6.3 Days 61-90: Build the hunt product.

Define the hunt cell's deliverable cadence: monthly novel rule deployment count, quarterly adversary playbook documentation, quarterly board readout. First hunt readout to Risk Committee at day 90.

Phase	Deliverable	Owner	Board Touchpoint
Days 1-30	Tiered alert library + playbook drafts	CISO + SOC Lead	Charter
Days 31-60	50% T-A migration + co-pilot live	CISO + CTO	Update
Days 61-90	Hunt cell stood up + first readout	CISO	Risk Committee
Day 90+	Quarterly Hunt Readout	SOC Lead via CISO	Standing

BOARD RECOMMENDATIONS

Decisions the board must take this quarter.

#	Decision	Owner	Evidence Required
R01	Adopt the three-tier SOC model with signed playbooks for every T-A alert.	CISO	Tiered library + playbook register
R02	Treat hunt as a board-reported product, not an SOC byproduct.	Board	Quarterly hunt readout
R03	Mandate tool entry-criteria: tuning + playbook + tier before deployment.	CISO	Procurement gate update
R04	Track analyst attrition as a Tier-1 board metric alongside detection time.	RemCo	People metric pack
R05	Sign the quarterly SOC Effectiveness Attestation personally.	CISO	Sign-off + minutes

A SOC re-architected on three tiers produces hunt where there was once drowning. The board funds the hunt, not the queue.

REGULATORY CROSS-WALK

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

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 Drowning in Noise.

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 Drowning in Noise.

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	Drowning in Noise operational dashboard	CISO function	Risk Committee minute
Quarterly	Drowning in Noise 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 Drowning in Noise.

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	Drowning in Noise 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.

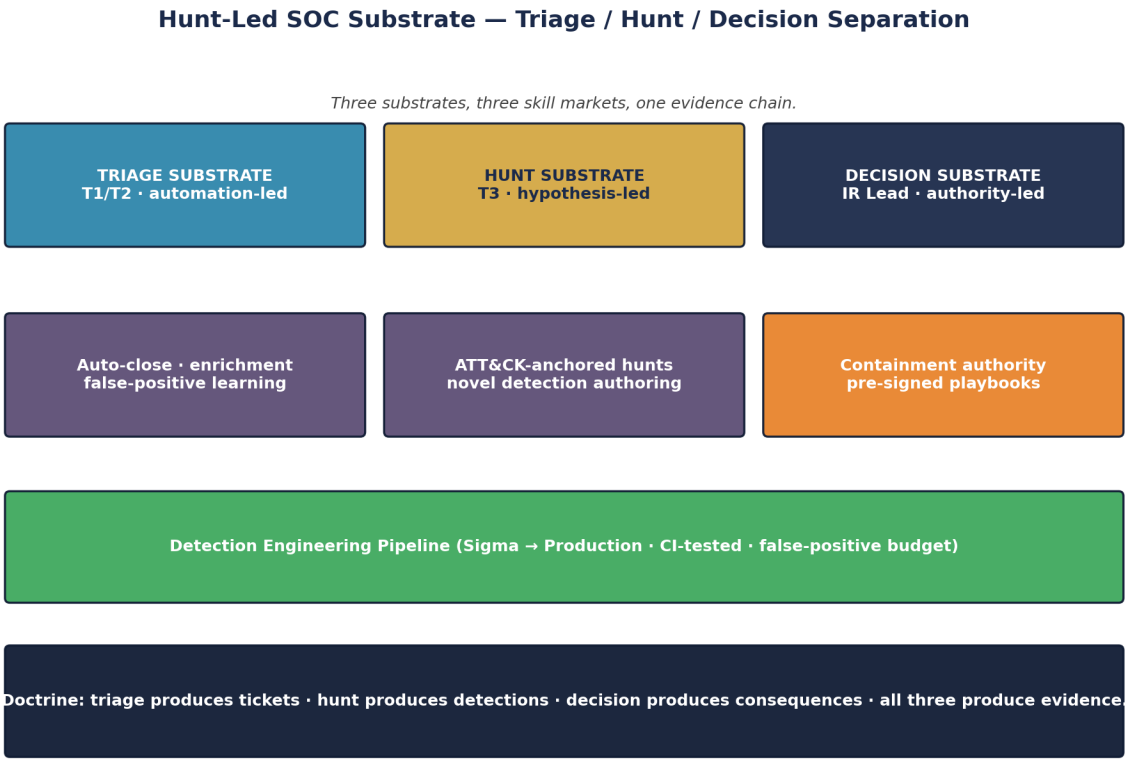

Figure A.P05. *Reference architecture for the doctrine in this paper. Colour coding: red denotes adversary or threat surface; teal denotes telemetry and detection; gold denotes classification and arbitration; navy denotes governance and decision authority; orange denotes human-in-loop; green denotes 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 — SOC Substrate Separation

```yaml
# soc_substrates.yaml — three substrates, one evidence chain
substrates:
  triage:
    workload: known_signature_alerts
    automation_target: 0.85   # 85% auto-closed
    sla_minutes: 30
    evidence: tickets/triage.parquet

  hunt:
    workload: hypothesis_based_threat_hunting
    framework: mitre_attack
    output: new_detection_rules
    cadence: continuous
    evidence: hunts/<hypothesis_id>.md

  decision:
    workload: incident_response
    authority: pre_signed_decision_rights
    pre_signed_actions:
      - isolate_host
      - revoke_credential
      - notify_regulator
      - communicate_externally
    evidence: decisions/<incident_id>.signed
```

Python — Detection Engineering CI

```python
# detection_ci.py — run on every Sigma rule before production
import yaml, subprocess

def validate_rule(path: str) -> dict:
    rule = yaml.safe_load(open(path))
    # 1. Lint
    assert rule.get('logsource'), 'no logsource'
    assert rule.get('detection'), 'no detection block'
    # 2. False-positive budget
    fp = subprocess.check_output(['fp_simulate', path]).decode()
    fp_rate = float(fp.strip())
    assert fp_rate < 0.005, f'FP rate {fp_rate} exceeds budget'
    # 3. ATT&CK coverage check
    assert rule.get('tags', {}).get('attack'), 'no ATT&CK tag'
    return {'rule': path, 'fp_rate': fp_rate, 'status': 'pass'}
```

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 · FRAMEWORK

Hunt-Led SOC Substrate™ — Definition, Falsifiability, Worked Calibration

Definition. A SOC operating model with three separated substrates — triage (automation-led), hunt (hypothesis-led), decision (authority-led) — each with its own skill market, evidence chain, and SLA.

Voice anchor. *Triage produces tickets. Hunt produces detections. Decision produces consequences.*

Aspect	Statement
Falsifiable claim	Hunt-Led SOC Substrate™ 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.

> ***"An alert queue is not a defence. It is a paper trail to a regulator."***

V3.0 · PRIMARY RESEARCH

Upadrasta Primary-Research Datasets — Cited In This Paper

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

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

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

V3.0 · MATURITYLADDER

Self-Service Maturity Scorecard — Where Is Your Institution?

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

Level	Description
1. Pre-Foundation	Single SOC tier. Analysts triage and respond simultaneously.
2. Foundation	Tier 1 / 2 / 3 model. No detection engineering function.
3. Operational	Detection engineering pipeline live. ATT&CK-mapped; hunts.
4. Institutional	Three substrates separated. Authority pre-signed.
5. Doctrine-Grade	Hunt outputs new detections weekly; FP budget < 0.5%.

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	Eight-week SOC Substrate 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 alert pathology, builds the substrate separation plan, designs the detection-engineering CI.. 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)	Microsoft Sentinel / Splunk / Elastic (SIEM substrate) · MITRE ATT&CK; (hunt anchoring) · External MSSP (24/7 triage substrate, retained)
Sector-First Reading	Critical National Infrastructure — NIS2 expects detection beyond signature.
Cyber-Insurance Position	Insurers credit institutions with an established detection-engineering function; pure SOC-as-MSP arrangements are increasingly under-rated.
M&A Cyber Due Diligence	Acquirer should ask: 'show me your last 10 detection rules deployed'. If response is 'we use vendor-default rules', it is a finding.
Litigation Defensibility	Negligence theories will probe whether the SOC was structured to detect the specific TTPs known to threaten the institution.
Board Sub-Committee Owner	Risk Committee + Audit Committee

V3.0 · N A V I G A T I O N

How To Read This Paper · Engagement Specialisms · ROI Envelope

How to read this paper.

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

Engagement Specialisms.

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

Indicative ROI envelope (this paper's doctrine).

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

V3.0 · CLOSING

Closing Doctrine — Paper-Specific

"An alert queue is not a defence. It is a paper trail to a regulator."

Hunt-Led SOC Substrate™

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 · M E T H O D

Methodology, Evidence Standards, and Sample Construction

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

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

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

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

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

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

TIER 1A · C I T A T I O N S

Primary Source and Citation Index

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

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

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 / MITRE
Detection engineering	Art. 10(1)	Art. 21(2)(b)	DE.CM-01	A.8.16	SYSC 13.7
Hypothesis-led hunting	Art. 10(4)	Art. 21(2)(b)	DE.AE-04	A.8.16	MITRE ATT&CK
ATT&CK coverage	Art. 24	Art. 21(2)(f)	ID.RA-04	A.5.7	TIBER-EU
Substrate separation	Art. 5(2)	Art. 20(1)	GV.RR-02	A.5.2	SYSC 13.6
Authority-led decision	Art. 11(2)	Art. 21(2)(c)	RS.MA-01	A.5.24	SYSC 13.8
False-positive budget	Art. 10(5)	Art. 21(2)(b)	DE.AE-03	A.8.16	SYSC 13.7
Sigma rule pipeline	Art. 10(1)	Art. 21(2)(b)	DE.CM-09	A.8.16	SYSC 13.7

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
Hunt-Led SOC Substrate™	Author framework: separation of triage / hunt / decision into three substrates with distinct skill markets.
Detection Engineering	Discipline of authoring, testing, deploying, and retiring detection rules with CI / CD discipline.
Sigma Rule	Open detection-rule format for SIEM-agnostic content sharing.
False-Positive Budget	Tolerance threshold for false-positive rate of a detection rule before retirement; institutional rules typically < 0.5%.
Hypothesis-Led Hunting	Threat hunting initiated from a documented hypothesis derived from threat intel or environment knowledge.
MITRE ATT&CK;	Knowledge base of adversary tactics and techniques; primary detection-coverage anchor.
Tier-3 Hunter	Senior SOC role focused on hypothesis-led hunting and detection authoring, distinct from triage and response substrates.

TIER 1A · S C O P E

Limitations, Scope, and Defensibility Caveats

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

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

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

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

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

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

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

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

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

THE CLOSING DOCTRINE

The doctrine in one line.

The drowning SOC is not a labour shortage; it is an architectural error. The doctrine separates decision from triage from hunt and assigns each to its proper substrate. Where the architecture is in place, the SOC produces evidence, novel detection, and recoverable defence. Where it is absent, the SOC produces ticket counts, analyst attrition, and the next 91-day dwell-time disclosure.

> ***"A drowning SOC reports activity. A hunting SOC reports adversary behaviour. The board pays for the second; it should never settle for the first."***

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 drowning SOC reports activity. A hunting SOC reports adversary behaviour. The board pays for the second; it should never settle for the first."

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