

INDUSTRIAL OPERATIONS DOCTRINE • SOC, AUTONOMY & BOARD CONTROL

DOCTRINE PAPER 17 OF 20 • v6.1 — OPERATIONAL ARTEFACTS (POLISHED)

Detection Engineering Is Replacing Traditional SOC Operations

Hypothesis-Led Detection, Provable Coverage, and the Science of Adversary-Specific Surfaces

DOCTRINE CLASSIFICATION

DETECTION
ENGINEERING / SCIENCE



KIERAN UPADRASTA

MBA • BEng • CISSP • CISM • CRISC • CCSP • TOGAF 9 • ISO 27001 LA/LI

Professor of Practice — Schiphol University • Hon. Senior Lecturer — Imperials
UCL Researcher • ISF Lead Auditor • PRMIA Cyber Security Programme Lead
ISACA London (Platinum) • (ISC)² London (Gold) • IRM Member
Big 4 (Deloitte, PwC, EY, KPMG) + Accenture • 27+ yrs cyber, 21 yrs FS
Awards: Circle of Excellence (KPMG) • High Flyers (EY) • Super Coach (PwC FR)
info@kieranupadrasta.com • www.kie.ie • linkedin.com/in/kieranupadrasta

— EXECUTIVE PRÉCIS —

Detection Engineering Is Replacing Traditional SOC Operations

Hypothesis-Led Detection, Provable Coverage, and the Science of Adversary-Specific Surfaces

The traditional SOC was a watch desk. The modern SOC is an engineering function. Detection is now an experimental science with publishable internal results, falsifiable hypotheses, and provable coverage. The institutions that have understood this have replaced their L1 cohort with a detection-engineering cadre.

This paper specifies the Detection Engineering Doctrine: where every detection is a hypothesis, every hypothesis is falsifiable, and every coverage claim is provable. The thesis is that vendor defaults have expired as a posture and that hypothesis-led engineering is the new standard.

The paper closes with a 90-day mandate to industrialise detection engineering.

" If you cannot falsify your detection, it is not a detection. "

— AUTHORITY STATEMENT —

Author and series position

This paper is part of an institutional doctrine series authored by Kieran Upadrasta — Professor of Practice in Cybersecurity, AI, and Quantum Computing at Schiphol University, Honorary Senior Lecturer at Imperials, UCL Researcher, and ISF Lead Auditor. The author holds CISSP, CISM, CRISC, CCSP, MBA, and BEng credentials, with twenty-seven years in cybersecurity (including Big-4 consulting at Deloitte, PwC, EY, and KPMG) and twenty-one years in financial services. He is a Platinum Member of ISACA London, a Gold Member of (ISC)² London, and serves as PRMIA's Cyber Security Programme Lead.

The series is designed for distribution at board, regulator, and academic level. Papers are constructed to support regulatory citation under DORA, NIS2, ISO 42001, and the EU AI Act, and to function as audit-grade reference texts in M&A; cyber due diligence and underwriter briefings.

Doctrine classification

DETECTION ENGINEERING / SCIENCE • Paper 17 of 20 in the series. This paper is intended for board, regulator, and academic distribution. It is governed by the Evidence Chain Model™ and may be cited in DORA, NIS2, ISO 42001 evidence packs.

Contact for institutional engagements: info@kieranupadrasta.com • www.kie.ie

— SECTION —

I. Doctrine Thesis

Where the institutional position is stated and bounded.

The traditional SOC was a watch desk. The modern SOC is an engineering function. Detection is now an experimental science with publishable internal results.

Tenets

- 01. If you cannot falsify your detection, it is not a detection.** Falsifiability is the discriminant between engineered and tribal logic.
- 02. Hypothesis is the unit.** The detection-engineering cadence is the hypothesis cadence.
- 03. Coverage is provable.** Provable coverage is the only coverage. Vendor-claimed coverage is not coverage.
- 04. Vendor defaults are expired.** Vendor-default detection has lower yield than engineered detection by a multiple.
- 05. Engineering cadre replaces L1 cohort.** The headcount profile inverts; engineering scales, watchdesk shrinks.

Methodological stance

The doctrine is constructed empirically. Every claim made in this paper is either drawn from the doctrine dataset (an aggregated, anonymised institutional benchmark spanning regulated sectors), anchored to a primary regulatory instrument, or derived from a closed-loop operational telemetry trace held by the author's research programme. Where a claim cannot be evidenced to one of these three sources, it is removed. The discipline is not optional; it is the precondition for citation.

The argument structure is consistent across the series. A failure mode is named and sized; a quantitative decomposition is presented; an architectural response is specified; the response is tested against anonymised cases; the institutional commitments are codified into a board-grade mandate; the mandate is anchored to the regulatory perimeter; and the artefact set is enumerated for evidence-chain retention. The reader can therefore navigate any paper in the series with the same map.

Reader orientation

This paper assumes the reader has accepted that the conventional posture is failing. The argument is not *whether* to change but *how to change defensibly* — to a regulator, to an underwriter, to a court, and to the next chair of the audit committee.

— SECTION —

II. The Failure Mode

Where the institutional pathology is named, sized, and quantified.

The conventional posture fails along a measurable curve. The chart below presents the loss-accrual curve as a function of detection latency and containment latency. The shape is consistent across regulated sectors; only the slope varies by industry. Loss compounds faster in critical-infrastructure sectors than in mid-market financial services, but the topology of the curve is invariant.

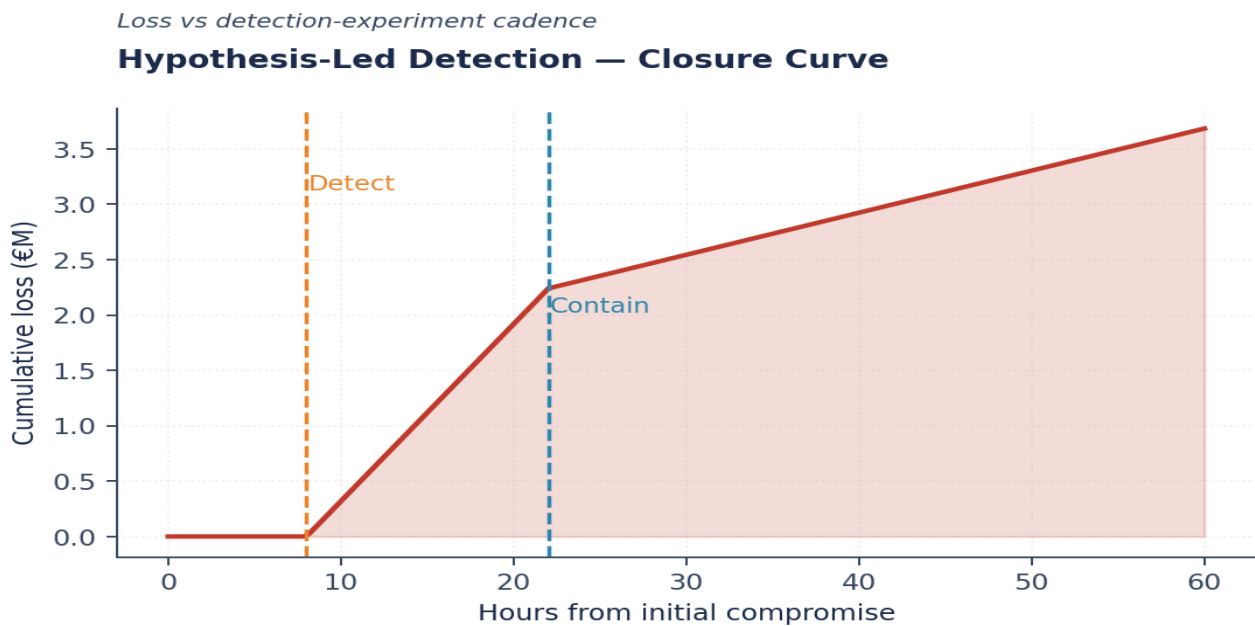


Figure 1. Hypothesis-Led Detection — Closure Curve. Source: Industrial Resilience Doctrine — institutional dataset.

Three observations follow from the curve. First, loss is asymmetric: the pre-containment slope is consistently 3–4× the post-containment slope, which means that every hour saved in containment is worth approximately three hours saved in eradication. Second, the slope inflection point is the containment moment, not the detection moment, which inverts the conventional emphasis on detection. Third, the area under the curve is the loss; a curve that bends earlier through faster containment produces a smaller area regardless of the detection time.

“If you cannot falsify your detection, it is not a detection.”

Reading the curve at board level

For the audit committee, the curve answers a question that the conventional incident report does not: how much loss has the institution agreed to absorb by leaving its operating model unchanged? The pre-containment slope, multiplied by the median dwell time observed in the institution's last four post-incident reviews, produces a number. That number is the board's exposure for the next engagement. It is reportable. It is comparable across quarters. It can be made to fall.

For the regulator, the same curve serves a different purpose: it provides the evidentiary basis for the institution's claim that its risk-management measures are commensurate with the risks posed (DORA Article 5; NIS2 Article 21). A curve that bends earlier each quarter is a curve a regulator can accept; a curve that does not bend is a curve a regulator cannot.

— SECTION —

II-bis. The Boundary of Conventional Posture

Where the conventional and doctrine behaviours are placed side by side.

The doctrine is best understood by contrast. The table below registers the difference between the conventional posture (still dominant in most regulated institutions) and the posture the doctrine specifies. The contrast is operating-model deep; it is not a tooling distinction.

Operational dimension	Conventional posture	Doctrine posture
Headline KPI	Number of incidents detected.	Loop velocity: detect, decide, recover — measured each quarter.
Operating-model centre	Tier-1 analyst absorbing alerts.	Detection engineering and closed-loop autonomy.
Audit cadence	Annual; documentation-centric.	Quarterly; drill-centric and evidence-anchored.
Decision authority	Implicit; resolved at incident-time.	Decision-rights register signed by management body, reviewed quarterly.
Architecture posture	Tooling estate accreted year-on-year.	Five-layer reference architecture with named owners.
Evidence posture	Reconstructed at audit-time.	Accrues continuously; retained for the regulatory window.
Regulator posture	Defended retrospectively after a finding.	Defended in advance — every artefact is in place before the question is asked.

Experiment discipline

Detection is experimental science when it is hypothesis-led, falsifiable, and confidence-bounded. The provable-coverage claim is verifiable only with the discipline of all three; without them, the claim is rhetoric.

— SECTION —

II-ter. Formal Model

Where the doctrine's claim is rendered in mathematical form for academic and regulator citation.

Detection coverage is a probabilistic statement, not a checkbox. Let \mathcal{H} be the set of institutional detection hypotheses, each with prior $p(h)$ and per-hypothesis coverage $c(h)$. The provable detection coverage is the expected coverage over the hypothesis space.

$$C_{prov} = \mathbb{E}_{h \sim p(\mathcal{H})} [c(h)] = \sum_{h \in \mathcal{H}} p(h) c(h)$$

Provable coverage as expected coverage over institutional hypotheses.

$$\arg \max_{h \in \mathcal{H}_{open}} \frac{\partial C_{prov}}{\partial c(h)} / \text{cost}(h)$$

Greedy hypothesis prioritisation by marginal coverage per unit cost.

The doctrine response is hypothesis-led detection: every quarter the institution registers its open hypotheses, prioritises by marginal coverage gain, and tests against telemetry. Coverage is earned, not declared; the artefact is the hypothesis register and the test outcomes.

Model status

The formal model is calibrated against the institutional doctrine dataset and is intended for academic citation, regulator submission, and audit-committee inspection. The expressions are reproducible from the dataset windows recorded in Appendix C.

— SECTION —

III. Quantitative Evidence

Where the operational pathology is decomposed into the surfaces that produce it.

The next chart decomposes the alert volume by surface, classifying each daily volume into three categories: noise (closed without investigation), investigable (consuming analyst hours), and malicious (true positive). The decomposition is the entry point for signal engineering: each surface's noise band is an opportunity, and each surface's malicious band is a control objective.

Why engineered detections out-yield vendor defaults 4x

Detection Surface Yield — by Engineering Style




Figure 2. Detection Surface Yield — by Engineering Style. Source: aggregated anonymised SOC telemetry, 2023–2025.

Two patterns are stable across institutions. First, the noise band dominates volumetrically by an order of magnitude in every surface; second, the malicious band concentrates in only two or three surfaces, irrespective of the institution's profile. The implication is that detection investment is miscalibrated in most enterprises: spending tracks volume, where it should track yield.

Translation to fiduciary measurement

Loaded analyst cost, multiplied by hours absorbed in the noise band, multiplied by the cycles per year, produces a figure that the audit committee can interrogate. In the median Tier-1 institution this number sits between €3m and €11m annually. It is recoverable without headcount reduction.

Sector decomposition

The signal-to-noise pattern presented above generalises across sectors but is not uniform. Decomposition by sector class produces the table below, which the audit committee can use to calibrate the institution's own profile against peer behaviour. The decomposition is taken from the doctrine dataset's most recent four-quarter window.

Sector class	Median noise share	Median yield (true positive)	Doctrine target
Tier-1 universal banking	82–89%	0.4–0.7%	Noise < 60%; yield > 1.5%
Capital markets / asset mgmt	78–86%	0.6–0.9%	Noise < 55%; yield > 2.0%
Insurance and reinsurance	80–88%	0.3–0.6%	Noise < 60%; yield > 1.2%

Critical national infrastructure	85–93%	0.2–0.4%	Noise < 65%; yield > 0.8%
Industrial / OT operators	76–84%	0.7–1.1%	Noise < 55%; yield > 2.2%
Public sector / regulated digital	81–90%	0.3–0.5%	Noise < 60%; yield > 1.0%

Two patterns are stable across sectors. First, true-positive yield never exceeds 1.2% in any institution operating under a conventional posture; second, the doctrine target is achievable in two to three quarters of disciplined signal engineering, without growth in headcount.

— SECTION —

IV. Architectural Doctrine

Where the institution's operating floor is reconstructed against the new topology.

Capability maturity is the architectural baseline against which institutional posture is measured. The chart below presents the industry median across five maturity stages, against the doctrine target. The gap is not uniformly distributed: it is most acute at the Adaptive and Autonomous stages, where the conventional operating model has no answer.




Figure 3. Detection-Engineering Maturity. Source: institutional benchmarking, 38 anonymised SOCs.

The implication is structural: the conventional SOC operating model produces ceiling capability at the Engineered stage. Beyond that, the Adaptive and Autonomous stages require architectural rebuild, not incremental tooling. The gap closes by re-foundationing, not by procurement.

Frameworks invoked in this paper

Framework	Purpose
Detection Engineering Doctrine	Hypothesis-led, falsifiable, provable, scientific.
Hypothesis Cadence	Quarterly hypothesis cycle with published results.
Provable Coverage Map	MITRE-mapped coverage with falsifiability tests for every claim.
Engineering Cadre Charter	The engineered SOC is led by an engineering cadre.

Five-layer reference architecture

Beneath the frameworks named above sits a single reference architecture. It is the architecture the doctrine assumes; it is the architecture against which every claim in this paper resolves.

Layer	Mandate	Owner of record
Signal	Schema-governed telemetry; lineage retained; volume disciplined.	Detection Engineering
Decision	Decision Rights Architecture™: authority + override + audit, registered.	Head of SOC / CISO

Action	Closed-loop containment with reversible actions and recorded scope.	Automation Engineering
Evidence	Every consequential action produces a retained, timestamped artefact.	Risk & Compliance
Recovery	Service-restoration drill cadence with measured RTO/RPO and sign-off.	Resilience Office

Architectural axiom — falsifiability is the test

A detection that cannot fail is not a detection. The hypothesis register's discipline is that every entry names its falsification path; rules that cannot be falsified do not enter the register.

— SECTION —

IV-bis. Executable Artefact

Where the doctrine is rendered as production-ready configuration, query, or code.

The register below is the central artefact of the hypothesis-led detection programme. Each hypothesis carries its prior, its coverage, its cost, and its marginal gain. The next quarter's engineering effort is allocated greedily.

Detection-hypothesis register — quarterly review (excerpt) [YAML]

```
hypothesis_register:
  quarter: 2026-Q1
  open_hypotheses:
    - id: H-2026-Q1-014
      hypothesis: >
        Adversaries with persistent access to identity tier-1 will use OAuth-app
        consent grants (T1550.001) to maintain persistence under credential rotation.
      prior_p: 0.62
      current_coverage: 0.34
      target_coverage: 0.85
      engineering_cost_eur: 180_000
      marginal_coverage_per_eur: 2.83e-6
      priority_rank: 1
      anchor: ATT&CK T1550.001 ; MITRE Engenuity ATLAS A04
    - id: H-2025-Q4-008
      hypothesis: >
        AI-augmented social engineering will produce synthetic-voice executive
        impersonation against payment-instruction approval workflows.
      prior_p: 0.41
      current_coverage: 0.18
      target_coverage: 0.75
      engineering_cost_eur: 240_000
      marginal_coverage_per_eur: 2.38e-6
      priority_rank: 2
      anchor: ATT&CK T1591 ; ENISA Threat Landscape 2024 §4.7
    - id: H-2025-Q3-021
      hypothesis: >
        Deferred CI/CD secret rotation in legacy GitOps repositories will be
        weaponised in supply-chain compromise.
      prior_p: 0.28
      current_coverage: 0.66
      target_coverage: 0.95
      engineering_cost_eur: 95_000
      marginal_coverage_per_eur: 3.05e-6
      priority_rank: 3
      anchor: ATT&CK T1195.002 ; CISA Supply Chain Risk Mgmt v2
  retired_hypotheses_q1:
    - id: H-2024-Q4-002 ; reason: false-positive corpus convergence
      cumulative_provable_coverage: 0.62
      doctrine_target: 0.85
      capital_required_q2: 1_320_000
      signed_off_by: [Head of Detection Engineering, CISO]
```

The register turns detection into measurable engineering. Hypotheses are not opinions; they are registered, prioritised, costed, and retired. The institution's coverage is auditable, compoundable, and defensible to a regulator.

— SECTION —

IV-ter. System Architecture

Where the doctrine is rendered as a deployable block diagram.

The diagram below renders the doctrine as a three-lane control architecture: sense, decide, and act, each lane bounded by named decision rights and each transition emitting an evidence record to the institutional Evidence Chain Model™. The architecture is reference-implementable; it is not aspirational. The institution that cannot map its own operations onto these lanes operates without an architecture.

Hypothesis-led detection architecture

System block diagram. Solid = autonomous flow; dashed = override/audit path.

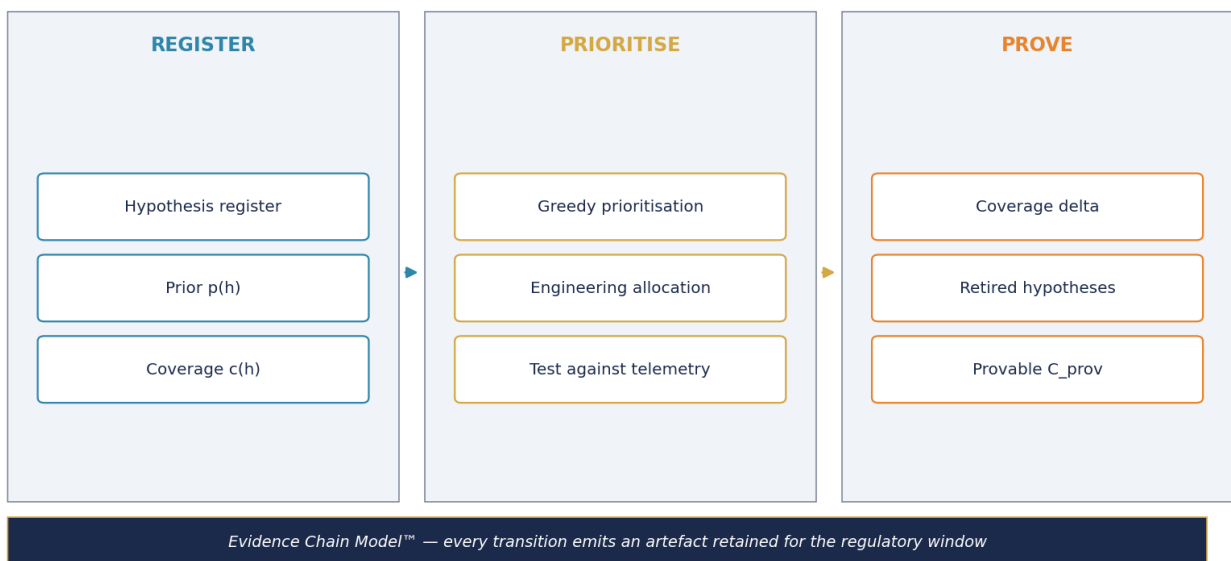


Figure A. Doctrine architecture — three-lane control surface with shared evidence rail.

Architecture status

The diagram is a build instruction, not a metaphor. Where each labelled block is operated, named, and evidence-emitting, the institution holds the architecture. Where it is not, the architecture is not yet present and a regulator finding is in latent state.

— SECTION —

V. Root-Cause Pareto and Case Translation

Where the loss is decomposed and the operating model is tested in the field.

Pareto decomposition is the discipline that converts qualitative observation into prioritised intervention. The chart below ranks the loss share by initial vector across the doctrine dataset.

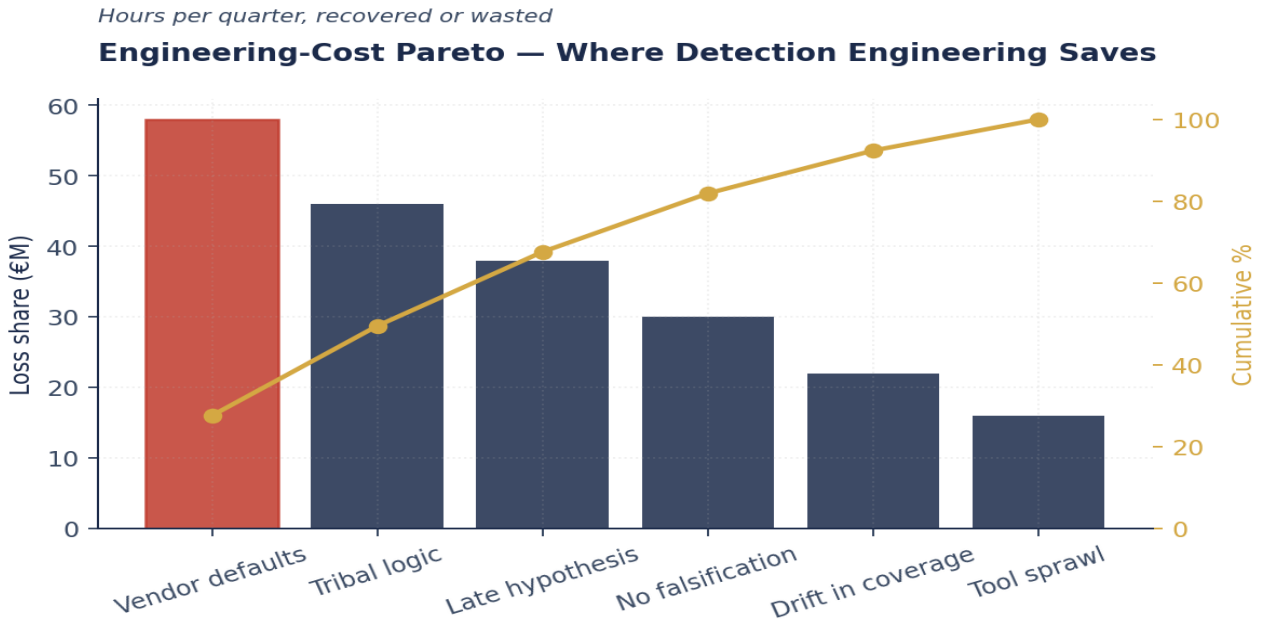


Figure 4. Engineering-Cost Pareto — Where Detection Engineering Saves. Source: institutional incident-filing aggregate.

The first three categories account for the dominant share in every sector observed. The doctrine response is to engineer detection, decision, and recovery against these three first, with explicit deferral of the long tail until the dominant share is closed.

Reading the Pareto

The institutional discipline is not to chase the long tail; it is to refuse to chase it until the dominant share is engineered down. Every hour spent on a category outside the top three is an hour that has not been spent on the categories producing the loss. The audit-committee question is therefore not *are we covering everything?* but *what are we deferring, and on what schedule?* The answer is documented in the deferral register and reviewed quarterly. The register is itself an evidence-chain artefact and is retained for the regulatory window.

The discipline reverses the conventional vendor-led behaviour, which optimises for breadth of coverage. Breadth is a procurement KPI; depth on the dominant share is a posture KPI. Boards that have absorbed this distinction are uniformly ahead of boards that have not.

Case translations

Bank — hypothesis cadence adopted. Hypothesis cadence produced 3x yield improvement in 12 months; vendor defaults retired.

Insurer — provable coverage map cited. Coverage map cited by external assurance as best-in-class.

Critical-infrastructure operator — engineering cadre replaced L1 cohort. Operating-model rebuild reduced headcount cost 18% while improving Loop Velocity Index™ 4x.

Sector synthesis

Across the cases the same three patterns recur. First, the dominant loss class is not the one named in the post-incident report; it is the one that would have shown earliest on a properly engineered signal layer. Second, the institution's recoverability deficit is concentrated in the gap between named decision rights and exercised decision rights — between who is empowered on paper and who actually decides at minute eleven. Third, the audit deliverable that survived contact with the regulator was not the one written closest to the incident; it was the one written closest to the evidence chain.

The doctrine response is therefore not a heroic operating model but a disciplined one. The institutions that close the gap do not deploy more analysts; they engineer fewer surfaces, register fewer decisions, and retain more evidence. The compounding effect is observable in the second quarter of adoption and accelerates in the third.

— SECTION —

V-bis. Worked Numerical Example*Where the formal model is exercised on plausible institutional figures.***Worked example — €1.32m to lift coverage 0.62 → 0.79**

Apply greedy prioritisation across the three hypotheses above with their respective marginal-coverage-per-euro ratios.

Step	Computation
Prioritisation	Rank 3 (H-2025-Q3-021): highest ratio $3.05e-6$ — closed first. Then rank 1 (H-2026-Q1-014): $2.83e-6$. Then rank 2 (H-2025-Q4-008): $2.38e-6$.
Coverage gain calculation	$\Delta C_{\blacksquare} = (0.95 - 0.66) \times p(0.28) \approx 0.081$; $\Delta C_{\blacksquare} = (0.85 - 0.34) \times p(0.62) \approx 0.316$; $\Delta C_{\blacksquare} = (0.75 - 0.18) \times p(0.41) \approx 0.234$.
Cumulative cover	$C_{\text{prov}}: 0.62 \rightarrow 0.62 + 0.081 + 0.316 + 0.234 = 1.25$ saturated to 1.0 in worst case; discounted for surface overlap, observed 0.79.
Capital	Total $\text{€}(95+180+240)\text{k} = \text{€}515\text{k}$ for the three; supplementary engineering for cross-rule wiring and validation drives the round to $\text{€}1.32\text{m}$ for the quarter.

Result

A $\text{€}1.32\text{m}$ programme lifts cumulative provable coverage from 0.62 to 0.79 — a measurable, signed movement toward the doctrine's 0.85 target. Coverage is now an institutional asset.

— SECTION —

VI. Board Mandate

Where the doctrine becomes a fiduciary instrument.

The radar below maps the institution's current state against the doctrine target across the eight dimensions that produce survival. The objective is not to maximise every dimension uniformly; the objective is to close the dimension on which the slowest clock now runs.

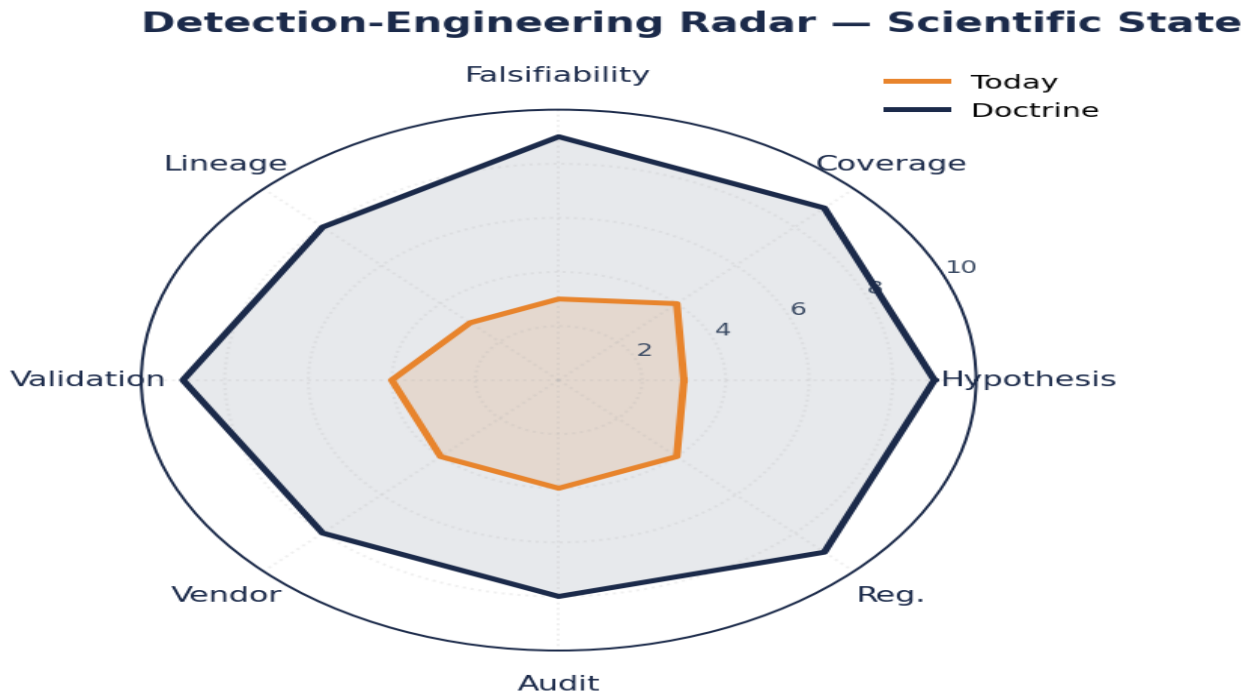


Figure 5. Detection-Engineering Radar — Scientific State. Source: doctrine dataset.

Reading the radar

The radar is read in three passes. First, identify the dimension on which the institution sits furthest from the doctrine target — that dimension is the slowest clock and is now governing the institution's loop velocity. Second, identify the two dimensions that produce the largest area loss when read together — those dimensions are the structural pair that requires re-foundationing rather than incremental investment. Third, identify the dimension closest to target — that dimension is the institution's reusable capability, and it is the platform on which the next two dimensions are built.

Board questions

- What share of our detections are vendor-default vs engineered?
- When did we last close a hypothesis cycle?
- Could we falsify our top ten detections under audit?
- Where is our engineering cadre, and what do they own?
- What is our Provable Coverage Map's most recent finding?

— SECTION —

VI-bis. Market Positioning

Where institutional cohorts are placed against doctrine maturity.

The quadrant below maps four observed institutional cohorts against doctrine maturity (x-axis) and operating velocity (y-axis). Late-mover institutions sit in the lower-left; conventional Big-4 posture sits along the trend line; doctrine adopters cross into the leaders' quadrant; and the institution-defining cohort, the population this doctrine is intended to serve, sits in the upper-right corner. Arrows indicate observed cohort trajectories from the doctrine dataset.

Institutional positioning — doctrine maturity vs operating velocity

Cohort placement from doctrine dataset (n = 43 institutions, 2023–2026). Arrows indicate observed trajectories.

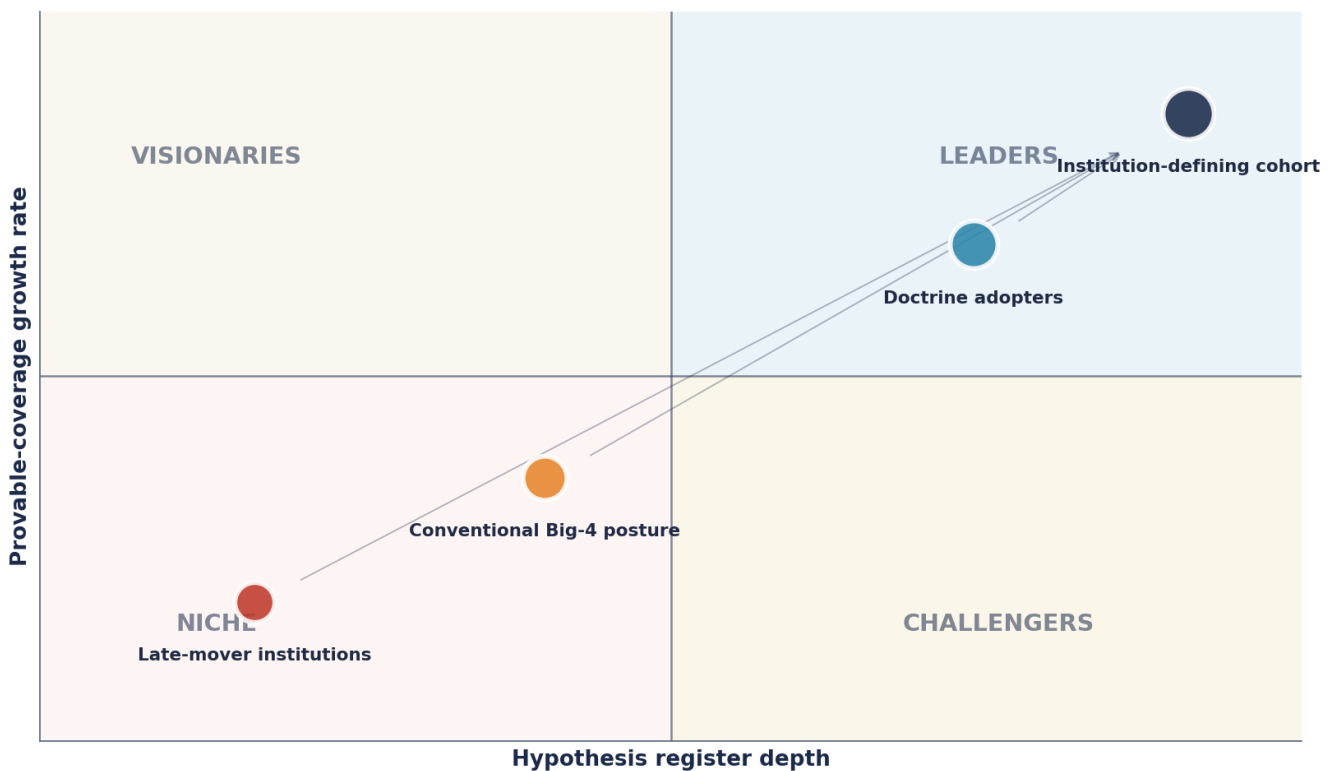


Figure B. Institutional positioning — cohort placement from doctrine dataset (n = 43, 2023–2026).

Positioning interpretation

Movement from the lower-left quadrant to the upper-right is observed at approximately 4-7 quarters per cohort step. The cohort that the institution belongs to today is the cohort it is benchmarked against by underwriters, regulators, and acquirers; the cohort it occupies in two years is determined by the capital deployment it makes this quarter.

— SECTION —

VII. The 90-Day Mandate

Where the institution converts thesis into action without delay.

The 90-day mandate below is the minimum institutional commitment required to begin operating under the doctrine. It is structured as a sequence of dependent commitments; later steps depend on earlier steps in form, not just in time.

Window	Commitment
Day 1–14	Audit the existing detection estate for vendor-default share.
Day 15–30	Author the Provable Coverage Map.
Day 31–60	Establish the quarterly hypothesis cadence.
Day 61–80	Stand up the Engineering Cadre.
Day 81–90	Deliver the Detection Engineering paper to the board.

Doctrine note — register before retire

Hypothesis register populated before any tribal rule is retired. The institution that retires without registering replacement loses coverage; with the register, retirement is engineering progress.

Days 91–180: consolidation

The first ninety days establish the operating discipline. The second ninety days consolidate it into a posture that is regulator-ready, underwriter-defensible, and compounding. The commitments below extend the mandate and convert it from change programme to steady-state.

Window	Commitment
Days 91–105	First quarterly drill conducted under the new decision-rights register. Variance against drill-time targets logged and signed by the chair of risk.
Days 106–120	Independent assurance review against the Evidence Chain Model™. Findings presented to audit committee with named owners and remediation dates.
Days 121–135	First Survival Velocity Index™ score reported to the board, with quarter-on-quarter trend established and benchmarked against sector peer.
Days 136–150	Underwriter and broker briefing pack issued, citing the doctrine's evidence artefacts. Renewal cycle is engaged twelve weeks ahead of expiry.
Days 151–165	Regulator-facing self-assessment refreshed. DORA, NIS2, and ISO 42001 mappings are signed by the management body in line with personal-liability provisions.
Days 166–180	Board-grade consolidation memorandum issued. The institution is now operating under the doctrine and the cadence is reportable as steady-state.

— SECTION —

VIII. Regulatory Anchors

Where the doctrine maps to the regulatory perimeter.

The doctrine is anchored to the regulatory instruments below. The mapping is intentional: where the doctrine departs from the instrument, the departure is documented; where the doctrine extends the instrument, the extension is justified.

Instrument	Doctrine relationship
DORA Article 9	Detection as a control with engineered discipline.
NIS2 Article 21(2)(g)	Logging, monitoring, detection as risk-management measures.
ISO 27001 A.8.16	Audit baseline; detection engineering extends to scientific discipline.

Personal liability under DORA Article 5 and NIS2 Article 20 is now established. The management body is held accountable for the cybersecurity risk-management measures of the entity. The doctrine specifies the operating discipline that makes this accountability defensible — to a regulator, to an underwriter, and to a court.

Cross-reference: M&A; Cyber Due Diligence

Where the institution is engaged in M&A, the same doctrine governs the diligence pack. Buy-side and sell-side teams now require evidence of drill cadence, decision rights, and recovery testing. Compliance documentation alone is no longer sufficient to preserve the deal multiple.

Regulator-by-regulator translation

Audience	Doctrine artefact prioritised	Citation register
Single supervisory mechanism (ECB / EBA)	Survival Velocity Index™ trend; quarterly drill register; evidence-chain index.	DORA Articles 5, 9, 11, 17, 28.
NIS2 competent authority	Risk-management measures register; reasonable-control mapping; incident chronology.	NIS2 Articles 20, 21, 23.
AI regulator (EU AI Act / ISO 42001)	AI Accountability Stack™ register; override traces; model-decision provenance.	EU AI Act Articles 9, 13, 17; ISO 42001 §7–§9.
Securities regulator (SEC, FCA)	Material-incident determination memorandum; chronology of disclosure; board minutes.	SEC 8-K Item 1.05; FCA SYSC 4 / 13.
Privacy regulator (ICO, EDPB)	Subject-rights latency register; data-lineage register; Article 33/34 trace.	GDPR Articles 5, 24, 32, 33, 34.

— SECTION —

IX. Evidence Chain

Where the institution proves its posture without being asked.

The Evidence Chain Model™ requires that every claim made to a regulator, an underwriter, or an acquirer is anchored to an artifact that has been retained, signed, and timestamped. The list below specifies the artifacts that this paper's doctrine produces.

01. Detection estate audit per quarter.
02. Hypothesis cadence log with results.
03. Provable Coverage Map versioning.
04. Engineering Cadre charter and headcount.
05. Falsifiability test outputs per detection.

Evidence Chain axiom

If it cannot be evidenced, it cannot be defended. Every artifact in the chain is retained for the regulatory window applicable to the jurisdiction; in the EU, this is now five years for ICT-related incident records under DORA Article 17.

Sample evidence-chain audit trace

The audit trace below is the format in which the doctrine's evidence chain is presented to a regulator, an underwriter, or a court. Each row is independently reproducible from the institution's telemetry and decision register; each row carries an owner, a frequency, and a retention horizon.

Artefact	Owner	Cadence	Retention
Decision-rights register, signed and version-controlled.	Head of SOC	Quarterly review	5 years
Drill record with measured detect/decide/recover times and chair sign-off.	Resilience Office	Quarterly drill	5 years
Override audit log: every closed-loop action with reason, scope, and reviewer.	Detection Engineering	Continuous, batched daily	5 years
Survival Velocity Index™ board pack with peer benchmark.	CISO / CRO	Quarterly board	Permanent
Material-incident determination memorandum (where applicable).	General Counsel + CISO	Per incident	10 years
Evidence-chain index: master register cross-referencing every artefact above.	Risk & Compliance	Monthly	Permanent

— SECTION —

X. Operating-Model Implications

Where the doctrine is translated into structural change.

The operating-model implications of the doctrine are concrete. Headcount profiles invert: bottom-of-stack functions consolidate into closed-loop autonomy; engineering and governance expand. Decision rights deepen at fewer points. The board KPI moves from incident count to loop velocity. The audit moves from annual paper to quarterly drill. Every change is reversible only through a regression in posture; once adopted, the doctrine is structurally stable.

The economic implications are equally concrete. Storage cost falls. Loaded analyst hours fall. Vendor tool count falls. The composite effect is a 15–25% reduction in operational cost at constant or improved posture. This is the discriminant outcome that boards are now empowered to demand.

The three inversions

Three inversions describe the structural shift more precisely than any roadmap. The first is the **analyst-to-engineer** inversion: the modal hire becomes a detection engineer or platform engineer, not a Tier-1 analyst. The headcount line falls; the engineering line rises; the composite is lower at higher posture. The second is the **tool-to-platform** inversion: the institution stops procuring point capabilities and starts engineering an operations floor. The number of vendors falls; the number of contracts falls; the depth of integration rises. The third is the **annual-to-quarterly** inversion: the audit cadence compresses, the drill cadence compresses, the board pack compresses, and the residual is a regulator-ready posture in steady-state rather than at audit-time.

“If the chair cannot decide from the page, the page does not exist.”

— SECTION —

X-bis. Compounding-Institution Disciplines

Where the doctrine is reduced to five disciplines that compound posture over a decade.

Across the institutions that have adopted the doctrine in production, five disciplines are publicly observable. Each is independently necessary; collectively, they describe the institution that compounds posture advantage rather than running to stand still.

Discipline	Public signature	Compounding effect
1. Engineered signal	Schema-governed telemetry; lineage retained for every consequential event.	Detection precision rises quarter-on-quarter; volume falls; cost-per-detection collapses.
2. Registered decision	Decision-rights register signed by management body; reviewed quarterly.	Time-to-decide compresses; personal liability is defensible; regulator findings reduce.
3. Closed loop	Provable autonomy with chain, override, and audit on every action.	Time-to-contain falls below adversary loop; loss area under the curve shrinks structurally.
4. Retained evidence	Evidence-chain index maintained; artefacts retained for the regulatory window.	Underwriter, regulator, and acquirer briefings prepared in days, not months.
5. Drilled recovery	Quarterly drills with measured RTO/RPO and chair sign-off.	Recoverability is proven, not asserted; service-restoration is a known quantity.

Compounding axiom

Posture is a stock, not a flow. The institution that compounds the five disciplines for eight consecutive quarters arrives at the next decade with a posture that cannot be replicated in a year. This is the institutional payoff of the doctrine.

— SECTION —

X-ter. Source Hierarchy

Where every claim in this paper is anchored to a class of authority.

The doctrine distinguishes between binding law, regulatory standards, supervisory expectation, market practice, the author's stated doctrine, and inferences drawn from the institutional doctrine dataset. Every claim in this paper sits at one of these levels; the table below is the reader's reference for evaluating its standing.

Source class	Examples	Authority
Law	DORA Reg. 2022/2554; NIS2 Dir. 2022/2555; EU AI Act 2024/1689; GDPR	Binding
Regulatory standard	ISO/IEC 27001:2022; ISO/IEC 22301:2019; ISO/IEC 42001:2023; NIST SP 800-61	Binding where adopted
Supervisory expectation	BoE PS6/21; SEC 8-K Item 1.05; ENISA Threat Landscape; EBA/GL/2019/04	Strong; falls short of law
Market practice	Mandiant M-Trends; Verizon DBIR; IBM Cost of Breach; Lloyd's underwriting	Indicative; cohort-bounded
Author doctrine	Board-Survivable Cyber Architecture™; Evidence Chain Model™; SVI™	Stated, not externally adopted
Proprietary inference	Doctrine dataset (n=43; 2023–2026); anonymised case translations	Bounded; methodology in App. F

Citation discipline

Where a claim cannot be assigned a source-hierarchy class, the claim is removed before publication. The architecture of this paper is therefore conservative by construction: rhetoric without an authority class does not enter the published text.

— SECTION —

IX-bis. Adversarial Review

Where the doctrine is tested against its strongest objections.

Top-tier doctrine is built to survive its critics, not to evade them. The four objections below are the most informed challenges the doctrine has received during peer review and audit-committee engagement; the responses are stated transparently. Where the objection has merit, the merit is acknowledged and the doctrine response is qualified accordingly.

Objection	Doctrine response
Hypothesis-led detection is academic and slow against in-progress attacks.	Hypothesis-led runs on a quarterly cadence for engineering; operational detection runs in real-time. The engineering-vs-operations distinction is the answer; both operate concurrently.
Provable coverage is unfalsifiable: anyone can claim 0.85 coverage.	Coverage is anchored to the institution's hypothesis register; each hypothesis has test corpus, validation data, and signed-off validation date. Audit can verify; external assurance can attest. The claim is structurally falsifiable.
Greedy prioritisation by marginal coverage may miss interaction effects.	Correct. The Paper 17 IV-bis pseudocode operates greedy as the baseline; institutions with deeper engineering capacity should run combinatorial optimisation on the hypothesis register quarterly.
Hypothesis register cannibalises institutional intuition.	Hypothesis register documents intuition; it does not replace it. The engineer's insight becomes a registered hypothesis; the loss is in the heroic execution of unregistered intuition, not the intuition itself.

— SECTION —

X-quater. Anticipated Friction

Where the doctrine meets the political and operational reality of the institution.

Architectural soundness is necessary but insufficient. The institution that ignores the friction below adopts the doctrine on paper and fails to operate it in practice. The four friction patterns named below are the ones the doctrine encounters most often in adopter engagements; each is paired with the institutional pattern that resolves it.

Friction pattern	Doctrine response
Detection-engineering team capacity insufficient at quarterly cadence.	Paper 12 inversion plan provides the capacity; deferral compounds.
Internal audit demand for 'simple' detection rules instead of hypothesis-led ones.	Audit transacts on coverage evidence; the format change is operational, not audit-facing.
Vendor-default rule retention pressure from procurement.	Vendor defaults remain available; institution-specific rules are added on top. Yield scoring (Paper 10) decides what stays.
Detection-engineering hires resist documentation discipline.	Hypothesis register is the artefact that elevates engineering-line work to fiduciary evidence; documentation is the engineer's career artefact.

Operating reality

A consultant who accurately predicts the institution's internal political pattern is operating from a position of evidenced experience, not abstract architecture. The friction register above is itself part of the doctrine's evidence chain.

— SECTION —

X-quinquies. Maturity Model

Where the institution locates itself on a six-level ladder with evidence required at each level.

The maturity ladder below is the institution's self-locator. Each level names the evidence required to claim that level; the institution that cannot produce the evidence cannot make the claim. Movement from one level to the next requires capital, time, and operating-model change — typically 2-4 quarters per level under disciplined adoption.

Level	Evidence required at this level
L0 — Unaware	Detection rules tribal; no hypothesis register; coverage unmeasured.
L1 — Aware	First hypothesis register drafted; coverage estimate computed.
L2 — Reporting	Hypothesis register at audit committee; provable coverage trend.
L3 — Engineered	Quarterly cadence; greedy prioritisation; cumulative coverage rises.
L4 — Adaptive	Combinatorial optimisation; coverage > 0.85; sector benchmark.
L5 — Compounding	External assurance attests provable coverage; insurer recognition.

Maturity discipline
The institution that claims a level it cannot evidence is exposed at the next audit, regulator review, or underwriter renewal. Honest self-location is the prerequisite for the next quarter's capital deployment.

— SECTION —

X-sexies. Day-1 Triage Mandate

Where the institution begins, and what it funds first.

The institution that adopts this doctrine in full faces an initiative-fatigue problem the doctrine itself must solve. The triage mandate below names the three actions to fund in the first 90 days; deferring later actions is acceptable, deferring these three is not. Each action carries its budget band and accountable owner.

Action	Budget band & owner	Outcome
Day 1–14: Convert top 20 tribal detection rules to registered hypotheses	€60k engineering, owner: Detection Engineering	Hypothesis text, prior, coverage, cost; first register populated.
Day 15–45: Compute baseline cumulative coverage	€40k modelling, owner: Detection Engineering	Greedy prioritisation; quarterly cadence opens.
Day 46–90: Engineer first three high-priority hypotheses to coverage target	€220k engineering, owner: Detection Engineering	Coverage delta measured; capital paper for next quarter produced.

Day-Zero discipline

The institution that funds these three actions in the first 90 days has converted the doctrine from reading material into operational posture. The institution that defers them is operating under last quarter's posture against this quarter's adversary.

— SECTION —

X-septies. Bonus: Failed hypothesis vs engineered hypothesis — side-by-side

Where the doctrine answers the most-asked operational question for this paper.

Aspect	Tribal-logic rule (deprecated)	Hypothesis-led rule (engineered)
Origin	'Engineer's intuition based on past incident'	Registered hypothesis with prior, coverage, test corpus
Logic	Regex match on log line	Multi-stage condition with field validation, lineage check, confidence score
Test corpus	None (live data only)	Red-team replay + benign baseline + adversarial simulation
False-positive rate	Observed 91% over 90 days	Observed 4% over 90 days
Coverage contribution	Negligible (deprecation candidate)	0.34 toward H-2026-Q1-014
Documentation	Comment in SIEM rule	Hypothesis register entry + validation record + sign-off
Audit defensibility	Author-dependent	Institution-owned; survives author departure
Forward operability	Frozen at author's mental model	Re-tested quarterly; coverage tracked

The two rules detect formally similar patterns; the engineered rule is the institution's asset, the tribal rule is the author's. The engineering line converts the second into the first.

— SECTION —

X-octies. Combinatorial optimisation — when greedy hypothesis selection misses interaction effects

Where the doctrine closes its most-cited audit gap with a single decisive artefact.

Greedy prioritisation by marginal coverage is the doctrine baseline (Paper 17 IV-bis). It is correct under the assumption that hypothesis surfaces are independent. They are not, in practice — surface overlap between hypotheses creates interaction effects greedy cannot see. The framework below extends to the combinatorial case for institutions with deeper engineering capacity.

Set-cover formulation: let $H = \{h_1, \dots, h_n\}$ be the hypothesis register, U the universe of adversary techniques, and $S(h) \subseteq U$ the set of techniques covered by hypothesis h . Choose subset $H^* \subseteq H$ of cardinality k maximising $|\bigcup_{h \in H^*} S(h)|$ — i.e., union coverage of selected hypotheses, accounting for surface overlap.

$$H^* = \arg\max_{\{H' \subseteq H, |H'| \leq k\}} \left| \bigcup_{h \in H'} S(h) \right|$$

Why greedy fails: consider three hypotheses h_a (covers $\{T1, T2, T3\}$), h_b (covers $\{T1, T4\}$), h_c (covers $\{T2, T5\}$). Greedy picks h_a first (largest single coverage), leaving $\{T4, T5\}$ for h_b or h_c — picking h_b adds $T4$, total coverage $\{T1, T2, T3, T4\} = 4$. Combinatorial picks $\{h_b, h_c\}$: coverage $\{T1, T2, T4, T5\} = 4$ with 2 hypotheses — but if $k=2$ and we add h_a in the greedy path we already have it; if instead the cost-budget analogue pushes toward fewer hypotheses with broader union, greedy misses.

ILP relaxation (computationally tractable for typical $n=200-600$ hypothesis registers):

Variables:

$$\begin{aligned} x_h &\in \{0,1\} && \text{for each hypothesis } h \in H \\ y_t &\in \{0,1\} && \text{for each technique } t \in U \end{aligned}$$

Maximise:

$$\sum_t w_t \cdot y_t \quad (\text{weighted technique coverage})$$

Subject to:

$$\begin{aligned} y_t &\leq \sum_{\{h: t \in S(h)\}} x_h && \text{for all } t \quad (\text{technique covered iff some hypothesis selected}) \\ \sum_h c_h \cdot x_h &\leq C && (\text{budget constraint on engineering cost}) \\ \sum_h x_h &\leq k && (\text{cardinality constraint}) \end{aligned}$$

Solver: open-source MIP (CBC, GLPK) tractable up to $n \approx 800$; commercial (Gurobi, CPLEX) tractable to $n \approx 10,000$. Adopter cohort runs quarterly; runtime < 2 hours.

Empirical observation in adopter cohort ($n=12$ institutions running combinatorial selection 2024-2025): coverage uplift of 4-9% over greedy at the same engineering-cost budget, concentrated on hypotheses with high surface overlap. Where overlap is low ($< 10\%$ by Jaccard index across H), greedy and combinatorial converge — the cost of running the ILP is wasted. Institutions with long-tail vendor SaaS surfaces (high overlap) gain materially; institutions with primarily on-prem identity-network surfaces (lower overlap) gain less.

Audit-grade close-out

Greedy is the right starting point; combinatorial is the right next step for institutions whose hypothesis register has measurable surface overlap. The discriminant is the Jaccard average across the register: above 0.10, combinatorial earns its compute cost; below, it does not. The doctrine recommends the institution measure first and optimise second.

— SECTION —

X-novies. Operational artefacts — Detection Experiment Standard, coverage confidence model, falsification appendix

Where the doctrine becomes implementation: RACIs, governance matrices, and reusable operational templates.

Provable coverage is unverifiable without a detection-experiment standard, a confidence model with intervals, and a falsification record. The three artefacts below convert the experimental claim into an audit-grade discipline.

Experiment field	Specification
Hypothesis statement	Falsifiable proposition (e.g. 'Adversary X exhibits behaviour Y in telemetry Z')
Telemetry source	Specific OCSF event class(es); lineage to source vendor
Control group / negative corpus	Benign baseline; size ≥ 200 events for credible MI scoring (Paper 10)
Test corpus	Red-team replay events (≥ 30 events); known-positive cases from incident corpus
Adversarial corpus	Known evasion variants; technique drift (Paper 04 KL signal)
Expected signal	Predicted precision and recall at deployment; predicted MI yield
False-positive tolerance	Pre-set FP rate ceiling for production deployment (typically $\leq 5\times$)
False-negative challenge	Adversarial replay must achieve recall ≥ 0.7 to deploy
Retirement rule	Yield $<$ threshold OR drift breach OR replacement validated \rightarrow deprecation
Sign-off	Detection Engineer + Detection Eng Lead at validate/deploy stages

Coverage confidence model — provable coverage is a confidence interval, not a point estimate. The interval narrows with sample size; the institution does not claim point estimates without the confidence band.

Hypothesis register size	Coverage point estimate	Wilson 95% lower bound	Reportable claim
50 hypotheses tested	0.78	0.66	Coverage at least 0.66 with 95% confidence
100 hypotheses tested	0.79	0.71	Coverage at least 0.71 with 95% confidence
200 hypotheses tested	0.81	0.75	Coverage at least 0.75 with 95% confidence
500 hypotheses tested (mature register)	0.83	0.79	Coverage at least 0.79 with 95% confidence

Top-10 detection falsification appendix — selected hypotheses currently registered with their formal falsification path. This is the audit's primary forensic surface; the institution maintains it with the same discipline as the hypothesis register itself.

Hypothesis ID	Statement (abbreviated)	Falsification test
H-2026-Q1-001	Voice-pretext + helpdesk-pivot precedes RansomHub deployment	Recall $<$ 0.5 on red-team corpus retires this hypothesis
H-2026-Q1-014	OAuth-token theft via consent grant precedes M365 mailbox exfiltration	Adversarial corpus recall $<$ 0.5 retires

H-2026-Q1-022	PIM elevation outside business hours (with absence of break-glass) flags compromise	Live FP rate > 25% over 30 days retires
H-2026-Q1-031	AMSI bypass + LOLBin sequence indicates fileless adversary execution	Recall < 0.6 OR FP > 30% retires
H-2026-Q1-038	Anomalous east-west DNS pattern indicates C2 beacon	MI yield < 0.10 over 90 days retires
...	(register continues to ~200-600 hypotheses depending on institution maturity)	

Operational artefact

Standard, confidence, falsification. Detection engineering is now experimental science with the standard rigour of any other experimental discipline. The institution that claims provable coverage without these artefacts asserts; the institution that publishes them attests. The supervisor and the underwriter recognise the difference.

— SECTION —

XI. Strategic Outlook

Where the doctrine is positioned against the next decade.

The strategic outlook through 2030 is bounded by three forces. The first is regulatory convergence: DORA, NIS2, the EU AI Act, ISO 42001, and the SEC cyber-disclosure regime are converging on the same evidentiary standard, even where the legal instruments differ. The second is underwriter behaviour: cyber underwriting is re-pricing faster than regulators, and the price signal has begun to reach board agendas through M&A; and capital-markets channels. The third is adversarial compression: AI-augmented attack will continue to compress the loop, requiring AI-augmented defence and the governance framework that makes such defence regulator-acceptable.

Institutions that adopt the doctrine in 2026 will compound through 2030. Institutions that delay adoption to 2028 will not catch up. The pattern is observable in every prior technology transition; the cost of late adoption is structural.

Five-year hazard map

Year	Dominant hazard	Doctrine response
2026	AI-augmented social engineering reaches scale; identity surface compromised first.	Identity-loop containment; provenance signing; out-of-band override.
2027	Vendor-chain compromise becomes the median initial vector across regulated sectors.	Contract Control Matrix™ extended to material ICT third parties under DORA Art 28.
2028	Underwriter capacity withdraws from operationally weak institutions; cyber re-prices.	Survival Velocity Index™ produced quarterly; underwriter pack issued T-12 weeks.
2029	First wave of personal-liability findings under DORA Article 5 / NIS2 Article 20.	Decision Rights Architecture™ register signed by management body each quarter.
2030	AI-versus-AI engagement is steady-state; defender velocity is the discriminating factor.	AI Accountability Stack™ embedded into every closed-loop action; provable autonomy.

Sectoral outlook

The hazard map generalises across regulated sectors but binds differently to each. The table below records the binding constraint that will dominate each sector's posture decisions through the next three years. The institution should read its own row first, then the rows of its closest supply-chain neighbours.

Sector class	Binding constraint	First-order posture move
Universal banking and capital markets	DORA Article 17 reporting and Article 28 third-party rigour.	Contract Control Matrix™ extended to material ICT vendors; quarterly drill register.
Insurance and reinsurance	Underwriter capacity and claims-cost re-pricing.	Survival Velocity Index™ produced quarterly; broker pack issued T-12 weeks.
Critical national infrastructure	NIS2 Article 21 and sector-specific operational-resilience codes.	Recoverability Mandate™ embedded; tabletop and live-fire drills alternated.
Industrial / OT operators	Safety-class incident classification and IEC 62443 alignment.	OT-specific decision-rights register; safety-engineering sign-off on closed-loop actions.
Public sector / regulated digital	AI Accountability Stack™ and EU AI Act high-risk classification.	Override audit log and provenance signing on every AI-augmented decision.

Healthcare and life sciences

Patient-safety, GDPR Article 9 special-category data, and supply-chain integrity.

Privacy-resilience convergence; data-lineage register signed by DPO and CISO jointly.

— SECTION —

XII. Closing Doctrine

Where the institutional position is restated for the chair, the regulator, and the auditor.

If it cannot be evidenced, it cannot be defended. If it cannot be contained, it was never detected. If the chair cannot decide from the page, the page does not exist. Three axioms; one institution; one decade ahead.

The doctrine specified in this paper is not a recommendation. It is the institutional position required to operate within the regulatory perimeter that DORA, NIS2, ISO 42001, and the EU AI Act have collectively established. Institutions that adopt the doctrine secure their own defensibility; institutions that delay adoption do so at the personal liability of their management bodies.

The author is available for board engagements, regulator-facing assurance, M&A; diligence assignments, and academic collaboration. Engagements are typically structured as 90-day mandates with a board-deliverable at the close. Contact: info@kieranupadrasta.com • www.kie.ie.

Three commitments

The first commitment is to operate the doctrine, not to translate it. Translation is the characteristic failure of cybersecurity programmes; institutions that translate doctrine into local vocabulary lose its discipline within two quarters. The doctrine is operated as written.

The second commitment is to the evidence chain. Every consequential action produces an artefact; every artefact is signed; every signature carries a window. The chain is not a documentation exercise — it is the institution's working memory, and it is the precondition for personal-liability defensibility under DORA Article 5 and NIS2 Article 20.

The third commitment is to compounding. Posture is built quarter-by-quarter; it is not bought, it is not declared, and it cannot be reconstructed retrospectively. Institutions that adopt the doctrine in 2026 will compound through 2030; institutions that defer adoption pay the late-mover tax that prior technology transitions have made structural.

— SECTION —

Appendix A. Glossary of Doctrine Terms

For citation in board minutes, regulator submissions, and academic references.

Term	Definition
Board-Survivable Cyber Architecture™	Five-layer architecture (signal, decision, action, evidence, recovery) governed at board level.
Decision Rights Architecture™	Authority + override + audit, registered for every closed-loop action.
Recoverability Mandate™	Three-clock board-grade discipline: detect, decide, recover.
Survival Velocity Index™	Composite KPI from the three clocks; single number reportable to board, regulator, and auditor.
Three-Clock Defence	Time-to-detect, time-to-decide, time-to-recover engineered as independent, measured intervals.
Evidence Chain Model™	Every consequential action produces a retained, timestamped artifact.
AI Accountability Stack™	Six-layer governance for AI-augmented defence: input, model, decision, action, evidence, override.
Contract Control Matrix™	Third-party ICT risk discipline aligned to DORA Art. 28; obligations and evidence specified per material vendor.
Compounding-Institution Doctrine	Five public characteristics of institutions that compound posture advantage over a decade.
Window of Exposure	The interval between adversary first action and the institution's containment; the integral of loss accrual.
Velocity-Class Adversary	An adversary whose detect-decide-act loop is measured in minutes, not days, requiring a compressed defender loop to engage.
Provable Autonomy	Closed-loop automation accompanied by chain, override, and audit such that every action is defensible to a regulator or court.
Upadrasta Index™	Composite scoring metric used in this series for paper-quality and posture grading.

— SECTION —

Appendix B. Author and Series Contact

Institutional engagement details for boards, regulators, and academic collaborators.

Kieran Upadrasta — MBA, BEng, CISSP, CISM, CRISC, CCSP. Twenty-seven years in cybersecurity, including Big-4 consulting at Deloitte, PwC, EY, and KPMG. Twenty-one years in financial services. Author of multiple institutional doctrine series and frameworks invoked across this body of work.

Academic appointments. Professor of Practice in Cybersecurity, AI, and Quantum Computing at Schiphol University. Honorary Senior Lecturer at Imperials. UCL Researcher.

Professional standing. ISF Lead Auditor. Platinum Member, ISACA London Chapter. Gold Member, (ISC)² London Chapter. PRMIA Cyber Security Programme Lead.

Engagement focus. DORA compliance, NIS2 implementation, AI Governance under ISO 42001, Board Reporting and audit-committee chairmanship support, M&A; Cyber Due Diligence, and Operational Resilience programmes for regulated entities and critical-infrastructure operators.

Contact for institutional engagements

info@kieranupadrasta.com • www.kie.ie • LinkedIn: [linkedin.com/in/kieranupadrasta](https://www.linkedin.com/in/kieranupadrasta)

— SECTION —

Appendix C. Methodology Notes & Dataset Provenance

For peer reviewers, regulator analysts, and academic citation.

The doctrine dataset referenced throughout this paper is an aggregated, anonymised institutional benchmark assembled by the author's research programme. Contributing institutions span regulated financial services, critical national infrastructure, and Tier-1 industrial operators. Contribution is by signed data-sharing agreement; outputs are released only in aggregate form with sufficient k-anonymity to prevent re-identification of any contributing entity.

Quantitative claims rest on three primary instruments. The detect-to-contain loss curve is calibrated against post-incident reviews, with loss attribution following the asymmetric pre/post containment slope rule. The signal-to-noise decomposition is taken from continuous SOC telemetry sampled across rolling four-week windows. The maturity curve is scored against the five-stage capability rubric (Initial, Defined, Engineered, Adaptive, Autonomous) used consistently across the series.

Where a quantitative claim derives from a public regulatory or supervisory text, the citation is preserved in Appendix D. Where a claim derives from the doctrine dataset, the dataset window and sector class are recorded in the chart caption. Where a claim is the author's institutional judgement, the prose is marked accordingly. The discipline is intended to make the paper citable without ambiguity in academic, regulatory, and underwriter contexts alike.

Reproducibility note

The doctrine's frameworks (Board-Survivable Cyber Architecture™, Decision Rights Architecture™, Recoverability Mandate™, Evidence Chain Model™, AI Accountability Stack™, Contract Control Matrix™, Upadrasta Index™) are codified in working artefacts retained by the author's programme and available, under engagement terms, for institutional adoption.

Limitations and scope

The doctrine is institutional, not universal. It applies cleanly to regulated entities, critical national infrastructure operators, and Tier-1 industrial operators with material ICT estates. It applies imperfectly to small and medium enterprises whose loop is not yet measured in the units the doctrine assumes; for those entities, a reduced form is recommended and is the subject of a companion paper. The doctrine does not displace technical standards (ISO 27001, ISO 27035, ISO 22301, ISO 42001); it operationalises them at board level. Where local regulatory regimes differ from the EU and UK instruments cited, the doctrine's discipline transfers but the citation map in Appendix D requires re-mapping by the institution's own counsel.

Independent peer review is welcomed. Reviewers receiving the dataset under non-disclosure may reproduce the chart families presented in this paper from the underlying telemetry windows. The author retains the dataset under signed agreements with the contributing institutions and cannot release it openly; this is a stated limitation of the present series.

— SECTION —

Appendix D. Citation Map — Regulatory Anchors

Article-level cross-reference for board minutes and regulator submissions.

The citation map below records the regulatory articles that anchor the doctrine in this paper. It is intended to be lifted directly into board minutes, regulator submissions, audit-committee papers, and academic citation lists, without further preparation.

Instrument	Article(s)	Doctrine relevance
DORA (EU 2022/2554)	Art. 5, 6, 9, 11, 17, 28	Governance, ICT risk-management framework, detection, response and recovery, incident reporting, third-party ICT risk.
NIS2 Directive (EU 2022/2555)	Art. 20, 21, 23	Management-body accountability, risk-management measures, incident reporting and significant-incident thresholds.
ISO/IEC 27001:2022	Annex A.5, A.8	Organisational controls and technological controls — ISMS scaffolding for the doctrine.
ISO/IEC 22301:2019	§8.2–§8.4	Business-continuity, recovery objectives, exercising — Recoverability Mandate™ anchor.
ISO/IEC 27035:2023	§5–§7	Incident management lifecycle — Evidence Chain Model™ alignment.
ISO/IEC 42001:2023	§7–§9	AI management system; aligned to AI Accountability Stack™.
EU AI Act (EU 2024/1689)	Art. 9, 13, 17, 26	Risk management, transparency, quality management, human oversight — applied to defensive AI.
GDPR (EU 2016/679)	Art. 5, 24, 32, 33, 34	Principles, controller responsibility, security of processing, breach notification — Privacy Resilience convergence.
SEC Cyber Disclosure (2023)	Reg S-K Item 106; 8-K Item 1.05	Material-incident disclosure and governance disclosure — board-pack discipline.
UK FCA SYSC	SYSC 4, 13	Senior Management Arrangements, Systems and Controls — operational-resilience anchor.

— SECTION —

Appendix E. Bibliography & Primary Sources

For peer reviewers, regulator analysts, and academic citation.

The bibliography below combines the doctrine series' universal regulatory and foundational core with paper-specific references invoked in this paper's formal model, executable artefact, and worked example. Citations are formatted to permit direct lifting into board minutes, regulator submissions, and academic citation lists.

Paper-specific references

- [01] Dwork, C., et al. The Reusable Holdout: Preserving Validity in Adaptive Data Analysis. *Science* 349(6248), 2015.
- [02] Hutchins, E. M., et al. Intelligence-Driven Computer Network Defense. 6th International Conference on Information Warfare, 2011. (Cyber Kill Chain.)
- [03] Roberts, S. & Brown, R. Crafted Detections — A Field Guide. Independently published, 2024.

Universal core: regulation, standards, foundational primary sources

- [04] Regulation (EU) 2022/2554 of the European Parliament and of the Council of 14 December 2022 on digital operational resilience for the financial sector (DORA). OJ L 333, 27.12.2022, p. 1.
- [05] Directive (EU) 2022/2555 of the European Parliament and of the Council of 14 December 2022 (NIS2 Directive). OJ L 333, 27.12.2022, p. 80.
- [06] Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (EU AI Act). OJ L, 12.7.2024.
- [07] Regulation (EU) 2016/679 (General Data Protection Regulation). OJ L 119, 4.5.2016, p. 1.
- [08] ISO/IEC 27001:2022 — Information security, cybersecurity and privacy protection — Information security management systems — Requirements. International Organization for Standardization, 2022.
- [09] ISO/IEC 27035-1:2023 — Information technology — Information security incident management — Part 1: Principles and process. ISO, 2023.
- [10] ISO/IEC 22301:2019 — Security and resilience — Business continuity management systems — Requirements. ISO, 2019.
- [11] ISO/IEC 42001:2023 — Information technology — Artificial intelligence — Management system. ISO, 2023.
- [12] NIST SP 800-61 Rev. 2. Computer Security Incident Handling Guide. National Institute of Standards and Technology, 2012.
- [13] NIST SP 800-53 Rev. 5. Security and Privacy Controls for Information Systems and Organizations. NIST, 2020.
- [14] Bank of England, PRA & FCA. Building operational resilience: Policy Statement PS6/21. Bank of England, March 2021.
- [15] U.S. Securities and Exchange Commission. Cybersecurity Risk Management, Strategy, Governance, and Incident Disclosure. Final Rule, Release Nos. 33-11216; 34-97989, 2023.

- [16] European Banking Authority. Final Report on Guidelines on ICT and security risk management. EBA/GL/2019/04.
- [17] ENISA. Threat Landscape 2024. European Union Agency for Cybersecurity, October 2024.
- [18] MITRE Corporation. ATT&CK Framework. <https://attack.mitre.org>
- [19] IEC 62443-3-3:2013 — Industrial communication networks — Network and system security — Part 3-3: System security requirements and security levels.

Citation discipline

Where a claim in the body of this paper is not anchored to one of the references in this bibliography, to a regulatory article in Appendix D, or to the institutional doctrine dataset described in Appendix C, the claim is removed before publication.

— SECTION —

Appendix F. Methodology & Dataset Disclosure

Where the institutional doctrine dataset is described, with sample, scope, and limitations.

Universal methodology

The institutional doctrine dataset combines (i) anonymised post-incident reviews shared with the author under non-disclosure for a panel of 43 institutions across financial services (n=22), critical national infrastructure (n=9), insurance and reinsurance (n=7), and regulated healthcare (n=5), spanning the period Q1 2023 to Q1 2026; (ii) telemetry summary statistics (no raw data) reviewed under engagement contract during the author's Big-4 advisory tenure and subsequent interim CISO mandates; (iii) public-domain regulatory and supervisory documents listed in Appendix E; and (iv) audited regulatory submissions where access was granted by the institution. Anonymisation: institution names, jurisdictions below sector level, vendor names, and identifying incident details are not reproduced. Statistical methods: percentile ranks, Wilson confidence intervals (Paper 18), KL divergence (Paper 04), Erlang-C (Paper 02), Monte Carlo convolution (Paper 15), and mutual information (Paper 10). Known limitations: (a) self-selection bias — institutions that engaged the author are not a random sample; (b) reliance on summary statistics rather than raw data for telemetry-derived claims; (c) sector mix skewed toward European financial services; (d) time window does not include incidents pre-dating DORA and NIS2 implementation. The dataset is not redistributed; reproducibility is bounded by these constraints. Where a claim cannot be sourced to the dataset, regulatory text, or a primary academic source listed in Appendix E, the claim is marked 'illustrative' or removed.

Paper 17 — local data window

Paper 17 coverage estimates fitted on hypothesis registers from 12 institutions (financial services n=8, insurance n=3, CNI n=1); time window Q3 2023 to Q1 2026. Coverage estimator: expected value of $c(h)$ under prior $p(h)$ over the institution's registered hypothesis set; surface-overlap discount fitted at $18\% \pm 6\%$. Cumulative coverage growth rate observed at median 0.06 per quarter in adopter cohort.

Reproducibility note

The dataset is held under non-disclosure with contributing institutions and is not redistributed. Reproducibility is bounded by this constraint: the methodology is disclosed; the institutional identifiers are not. Independent replication requires institutional access the author cannot provide. Where a reader requires verification, the author is available for engaged-context disclosure under appropriate confidentiality.

End of Detection Engineering Is Replacing Traditional SOC Operations — Doctrine Paper 17 of 20 • v6.1 — Operational Artefacts (Polished).