

INDUSTRIAL OPERATIONS DOCTRINE • SOC, AUTONOMY & BOARD CONTROL

DOCTRINE PAPER 05 OF 20 • v6.1 — OPERATIONAL ARTEFACTS (POLISHED)

The Autonomous SOC Is No Longer Optional

Decision Rights, Provable Autonomy, and the Architecture of Trust at Machine Speed

DOCTRINE CLASSIFICATION

AUTONOMY / DECISION
RIGHTS



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 —

The Autonomous SOC Is No Longer Optional

Decision Rights, Provable Autonomy, and the Architecture of Trust at Machine Speed

Autonomy in the SOC is now a regulatory question, not a technological one. The decision is no longer whether to automate response, but how to evidence the autonomy such that a regulator, an auditor, and a board can all accept it.

This paper specifies Provable Autonomy — the discipline of closed-loop response that carries chain, override, and audit. It is the only autonomy that DORA, NIS2, and ISO 42001 will accept, and the only autonomy that will not collapse under cross-examination.

The paper closes with a 90-day mandate to convert the existing SOAR estate into provable autonomy and a board-survivable governance model.

" Provable autonomy is the only autonomy a regulator will accept. "

— 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

AUTONOMY / DECISION RIGHTS • Paper 05 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.

Autonomy in the SOC is now a regulatory question, not a technological one. The boundary between recommendation and authority is the boundary between liability and survival.

Tenets

01. Recommendation is no longer a defensible posture. When the adversary closes the loop in seconds, the defender that recommends within minutes has lost.

02. Provable autonomy is the only autonomy regulators accept. Every autonomous action must produce evidence of authority, evidence of override, and evidence of audit.

03. The override is the control. The boundary between autonomy and human authority is the override. Without a tested override, autonomy is unsupervised.

04. Audit is real-time. Post-hoc audit cannot supervise machine-speed action. The audit must be embedded in the action.

05. Decision Rights Architecture™ is the doctrine. Every closed-loop action carries authority, override, and audit. Without all three, autonomy is not governable.

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.

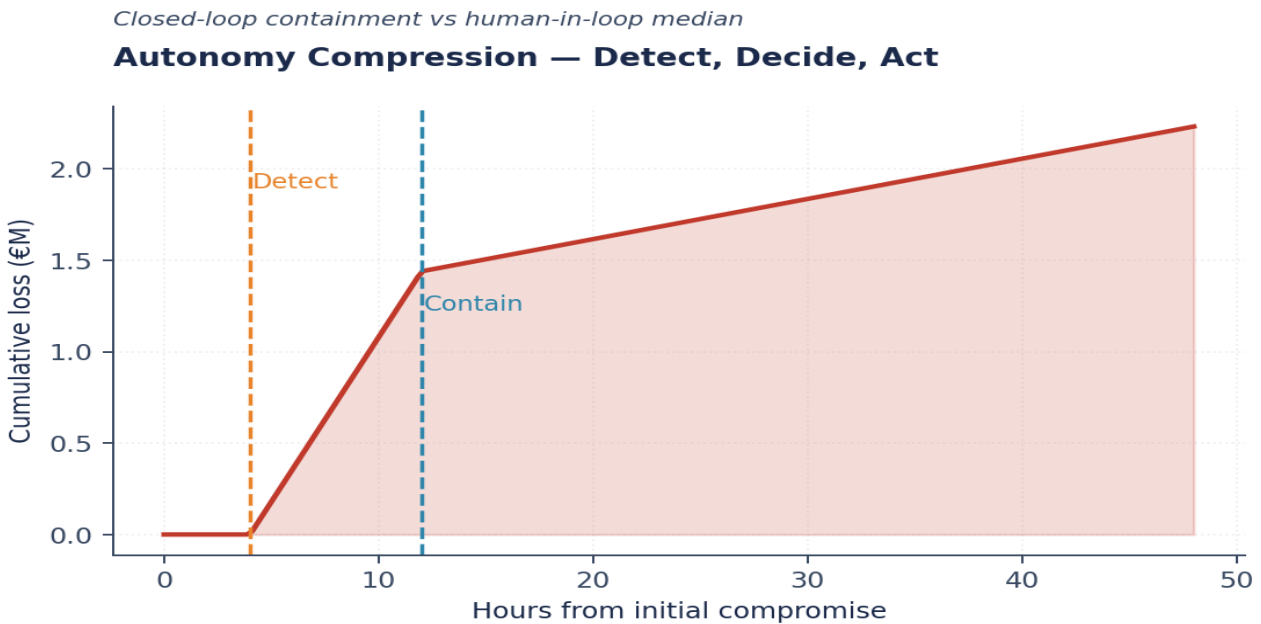


Figure 1. Autonomy Compression — Detect, Decide, Act. Source: Industrial Resilience Doctrine — institutional dataset.

Three observations follow from the curve. First, loss is asymmetric: the pre-containment slope is consistently 3–4x 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.

“Provable autonomy is the only autonomy a regulator will accept.”

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.

Autonomy discipline

Autonomy without authority is liability. The five tests below — chain, override, audit, reversibility, evidence — must be answerable for every closed-loop action before it operates, not after.

— SECTION —

II-ter. Formal Model

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

An autonomous SOC is a closed-loop control system. Let $r(t)$ be the threat reference signal (adversary action), $y(t)$ the institutional response, $e(t) = r(t) - y(t)$ the engagement error, and $K(s)$ the compensator (the SOC's decision-and-action stack). The closed-loop transfer function determines whether the system is stable, fast, and audit-defensible.

$$T(s) = \frac{K(s)G(s)}{1 + K(s)G(s)}, \quad e(t) = r(t) - y(t)$$

Closed-loop transfer with plant $G(s)$ (the institution) and compensator $K(s)$ (the SOC).

$$\text{Settling time } t_s \approx \frac{4}{\zeta\omega_n} \leq t_{regulator}^*$$

Doctrine constraint: settling time must clear the regulator's reporting horizon t^* (e.g. DORA's 4h).

The constraint binds. A SOC whose closed-loop settling time exceeds the regulator's horizon cannot produce a regulator-acceptable response by definition. The autonomous SOC is the engineering answer: the compensator $K(s)$ is closed in code, with provable autonomy — every action carries chain, override, and audit. Manual escalation is reserved for cases where override is invoked, not for normal operation.

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.

Bottom of stack collapses; top of stack expands

Hours Released by Autonomy — Where Effort Reallocates




Figure 2. Hours Released by Autonomy — Where Effort Reallocates. 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.

Provable autonomy carries control, evidence, and override

Autonomy Maturity — Reactive to Provable




Figure 3. Autonomy Maturity — Reactive to Provable. 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
Provable Autonomy Doctrine	Closed-loop response with chain, override, and audit; evidenceable to a regulator within the disclosure window.
Decision Rights Architecture™	Authority + Override + Audit, registered, tested, retained.
Override Drill	Quarterly test of every override path. If untested, the autonomy is suspended.
AI Accountability Stack™	Six layers: input, model, decision, action, evidence, override. Every layer auditable.

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 — reversibility is a design property

Reversibility is engineered before deployment, not asserted after misfire. The institution that designs the rollback path before it engineers the action operates closed-loop; the institution that engineers the action and assumes the rollback operates closed-loop in name only.

— SECTION —

IV-bis. Executable Artefact

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

The state machine below is the autonomous SOC's closed-loop control specification for endpoint containment. It is provably autonomous: every transition is logged, every action carries an evidence record, and override states are first-class.

Closed-loop containment — control state machine (excerpt) [YAML]

```
state_machine:
  name: closed-loop-endpoint-containment
  initial_state: monitoring
  evidence_class: control-loop-audit-record
  retention: 5y
  states:
    - id: monitoring
      transitions:
        - condition: signal.confidence >= 0.85 AND signal.surface == endpoint
          to: pre-action-validation
          emits: evidence.alert_triaged
    - id: pre-action-validation
      timeout_seconds: 30
      transitions:
        - condition: validator.consensus >= 2_of_3 AND override.flag == false
          to: contain
          emits: evidence.action_authorized
        - condition: override.flag == true OR validator.consensus < 2_of_3
          to: human-escalation
          emits: evidence.override_invoked
    - id: contain
      action: isolate_host
      transitions:
        - condition: action.success == true
          to: post-action-validation
          emits: evidence.action_executed
        - condition: action.success == false
          to: rollback
          emits: evidence.action_failed
    - id: post-action-validation
      transitions:
        - condition: telemetry.confirms_isolation == true
          to: monitoring
          emits: evidence.action_confirmed
        - condition: telemetry.does_not_confirm
          to: rollback
          emits: evidence.action_unconfirmed
    - id: rollback
      action: reverse_isolation
      transitions:
        - condition: rollback.success == true
          to: human-escalation
          emits: evidence.rollback_complete
    - id: human-escalation
      action: notify(role=SOC_Shift_Lead)
      sla_minutes: 5
      next_states: [contain, monitoring] # human decides
governance:
  decision_rights_register: DR-CONT-EP-001
  regulator_anchor:
    - DORA Art. 9(2)
    - DORA Art. 17(1)
    - EU AI Act Art. 14 (human oversight)
```

Every state transition produces an evidence record. The institution that operates this state machine produces, by construction, the audit trail required under DORA Article 17 and the human-oversight evidence required under EU AI Act Article 14. Provable autonomy is the only autonomy a regulator will now accept.

— 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.

Autonomous SOC closed-loop control architecture

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

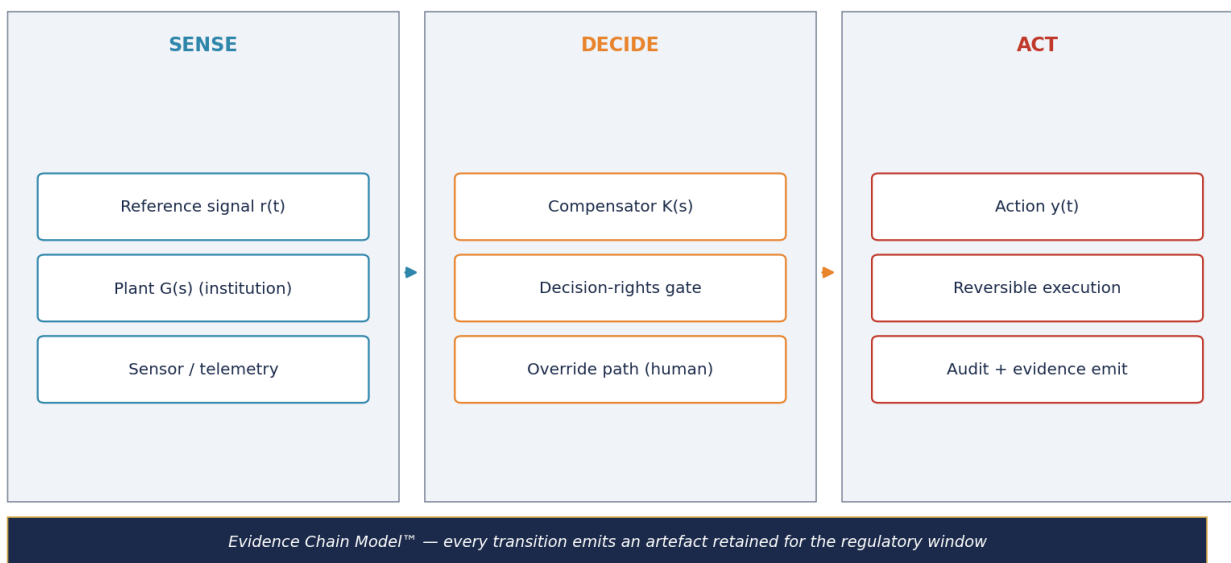


Figure A. Doctrine architecture — three-lane control surface with shared evidence rail.

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. Autonomy Pareto — Where Latency Is Recovered. 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 — DORA Article 5 closed-loop test. Provable Autonomy implementation produced a regulator-acceptable evidence chain; closed-loop response time 19 seconds; full audit trail retained.

Critical-infrastructure operator — autonomy under inspection. Regulator inspection of closed-loop response produced no findings; cited as a reference institution.

Insurer — override drill exposed weakness. Quarterly override drill revealed a path that had degraded silently. Doctrine repaired the path within seven days.

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 — settling time 4h → 11min**

An institution measures its current closed-loop characteristics and computes the gap to doctrine.

Step	Computation
Current parameters	Damping ratio $\zeta = 0.3$ (under-damped, oscillatory escalation); natural frequency $\omega_n = 0.0026$ rad/s (slow). Settling time $t_s \approx 4/(0.3 \times 0.0026) \approx 5,128$ s ≈ 1.42 hours.
Regulator constraint	DORA Article 17(3) initial-report horizon: 4 hours. Current t_s comfortably clears this horizon but only because the system rarely settles — most engagements escalate beyond 4h.
Doctrine target	$\zeta = 0.7$ (well-damped); $\omega_n = 0.011$ rad/s (faster, autonomous). $t_s' \approx 4/(0.7 \times 0.011) \approx 519$ s ≈ 8.7 minutes.
Posture implication	An 8.7-minute settling time clears the 4-hour regulator horizon by 27x; the institution operates with regulatory headroom, not regulatory drag.

Result

A measurable settling time (8.7 minutes) is the discriminant figure the audit committee uses to verify autonomous-SOC adoption. Where the figure cannot be produced, the autonomy claim is not made.

— 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. Autonomy Radar — Provable 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

- Which closed-loop actions in our environment have no named override?
- When did we last conduct an override drill, and what did it reveal?
- What is our chain-of-authority evidence for the last autonomous action that touched a customer?
- Could our regulator reproduce the audit of any autonomous action in the last 90 days?
- If our AI vendor disappeared tomorrow, do we still have provable autonomy?

— 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	Inventory every closed-loop action across SOAR, EDR, IDP, and cloud.
Day 15–30	For each, register authority, override, audit. Where any is absent, suspend the action.
Day 31–45	Conduct first override drill quarter. Capture failure modes.
Day 46–70	Stand up the AI Accountability Stack™ across the autonomous estate.
Day 71–90	Deliver the Provable Autonomy paper to the board. Three pages; one decision.

Doctrine note — sign before deploy

The Decision-Rights Register is signed before any closed-loop action enters production. Authorisation, override, evidence — all named on paper before a single autonomous action operates. The institution that deploys before signing operates outside the doctrine and outside the supervisor's tolerance.

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 5 (Governance)	Closed-loop autonomy must be evidenceable to the management body; the Decision Rights Architecture™ provides the structure.
NIS2 Article 21(2)(e)	Risk-management measures must include policies on the use of ICT, including any AI; provable autonomy converts policy into proof.
ISO 42001 Clause 9.2 (Internal audit)	AI-driven autonomy is subject to internal audit; the audit must be embedded, not retrospective.

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. Decision Rights Register for every closed-loop action.
02. Override drill outputs with quarterly cadence.
03. AI Accountability Stack™ audit trail.
04. Regulator-grade chain-of-authority document.
05. Vendor-action audit synced to internal evidence chain.

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
Autonomous containment will produce false-positive isolation and self-inflicted outage.	The doctrine response is the autonomy boundary matrix (Bonus below): only reversible actions are autonomous; all destructive or non-reversible actions require human authorisation. Reversibility is a first-class design property of the closed-loop state machine in Paper 05 IV-bis. The fear is real; the answer is engineered.
EU AI Act Article 14 (human oversight) prohibits closed-loop containment.	Article 14 mandates effective human oversight, which under the doctrine takes three forms: (a) override path always present and tested, (b) decision rights pre-signed by a named human (Decision-Rights Register), (c) post-action audit by human within 24h. Closed-loop in this configuration is not unattended autonomy; it is human-authored, human-overseen, audited execution.
Provable autonomy requires immutable audit logs the institution does not have.	Correct that immutable evidence chain is non-negotiable. The Evidence Chain Model™ (Paper 09 + Appendix B) defines the artefact set: signed action records, hash-chained audit trail, retention 5y under DORA Art. 17. Where these are absent, autonomy is not yet defensible and should not be deployed.
Closed-loop autonomy is acceptable only for endpoint isolation, not for identity or network.	False as a categorical claim. The doctrine permits closed-loop on any reversible action with sub-30-minute rollback window. Identity-tier-1 token revocation is reversible and is in scope; network segment cutover for a contained host is reversible and is in scope. Destructive actions (account deletion, key destruction) remain human-authorized.

— 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
Legal counsel concerns over the institution acting on autonomous evidence in court.	The Evidence Chain Model™ produces signed, hash-chained, retained-5y artefacts; litigation-grade by construction. The institution's general counsel is engaged in the Decision-Rights Register signoff (Paper 05 IV-bis).
CISO career-risk objection: 'autonomous mistake = my dismissal'.	Pre-action validation gate (2-of-3 validator quorum), reversibility, and audit cycle transfer the risk profile from blind autonomy to evidenced autonomy. The CISO's signed decision-rights register is the personal-liability defence under DORA Art. 5.
Operations leadership fear of automation 'taking over' a function they have run for 20 years.	Operations remains in the loop on override; closed-loop handles the deterministic 80%. Paper 09 cultural transition framework is the lift-out plan.
Vendor lock-in: the autonomy stack rests on a single SOAR vendor.	Schema sovereignty (Paper 11) and multi-platform decision-rights register (vendor-agnostic by construction) are the technical answer. The contract control matrix is the commercial answer.

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	All response is manual; no closed-loop actions; no decision-rights register.
L1 — Aware	First reversible closed-loop action identified; pilot in test environment.
L2 — Reporting	Endpoint isolation closed-loop in production; decision-rights register signed for endpoint class; settling-time measured.
L3 — Engineered	Identity tier-1 closed-loop deployed; override path drilled; evidence chain operational.
L4 — Adaptive	All reversible actions in closed-loop; ζ_{ω_n} meets regulator horizon by margin; automated post-action validation.
L5 — Compounding	External assurance attests provable autonomy; sector benchmark leader; insurer recognition; provable autonomy is the institutional posture.

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: Sign the first endpoint-class decision-rights register	€60k legal + ops, owner: CISO + General Counsel + CRO	Authority, override, evidence, retention all named. The artefact is the prerequisite.
Day 15–45: Deploy closed-loop endpoint isolation in production	€220k engineering + tooling, owner: Detection Engineering + Head of SOC	Rollback window 30 min; validator quorum 2-of-3; audit cycle 24h.
Day 46–90: Drill the override path under simulated false-positive	€80k drill, owner: Head of SOC + General Counsel	Drill evidence retained 5y; finding register opens; closure within 30d.

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: Autonomy boundary matrix — what closed-loop can and cannot do

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

Action class	Mode	Reversibility	Authority
Endpoint isolation (host quarantine)	Closed-loop	Auto-rollback ≤ 30 min	DR-CONT-EP-001 signed
Identity-tier-1 token revocation	Closed-loop	Auto-rollback ≤ 30 min	DR-CONT-ID-001 signed
Network segment cutover (contained host)	Closed-loop	Auto-rollback ≤ 30 min	DR-CONT-NW-001 signed
IAM role suspension (privileged)	Supervised auto	Reversible; audit ≤ 60 min	Human approval gate
Active-directory account deletion	Approval-gated	Non-reversible	CISO + Head of IAM
Key destruction / certificate revocation	Approval-gated	Non-reversible	CISO + General Counsel
Database wipe / backup destruction	Prohibited	Non-reversible	Never autonomous
Customer-facing service shutdown	Prohibited	High blast radius	Never autonomous

The matrix is signed by CISO, General Counsel, CRO, and COO; reviewed quarterly. Where the institution cannot sign a row in the right column, the row is not autonomous.

— SECTION —

X-octies. External assurance template — ISAE 3000 attestation for Provable Autonomy

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

Provable autonomy is unverifiable without external assurance. The template below is ISAE 3000 (Revised) Type 2 — the standard for assurance over non-financial subject matter — scoped to the institution's closed-loop containment architecture. The institution engages an independent practitioner (Big-4 audit firm or accredited cyber-assurance specialist) to issue the attestation.

Element	Specification
Subject matter	Operating effectiveness of closed-loop containment for endpoint isolation, identity-tier-1 token revocation, and network-segment cutover, including decision-rights register, override path, and evidence-chain emission.
Criteria	Suitable criteria: (i) ISO/IEC 42001:2023 clauses 6, 8, 9; (ii) DORA Reg. 2022/2554 Articles 5, 9, 11, 17; (iii) institution's own Decision-Rights Register; (iv) NIST AI RMF function-level controls.
Period	12 months ending [Q-end]; coverage of all closed-loop actions executed in the period (n events disclosed).
Practitioner	Independent firm with cyber-assurance accreditation; partner-led engagement; minimum 240 hours engagement scope for a Tier-1 institution.
Evidence sample	100% of closed-loop actions in period for design, statistical sample of at least 60 actions for operational testing; all override events tested.
Reporting	Reasonable assurance opinion (positive form). Limited assurance is insufficient for regulator and underwriter use under the doctrine.
Distribution	Audit committee, supervisor (DORA Art. 5 evidence), insurer renewal pack, M&A diligence pack.

Auditor sign-off statement (template wording):

INDEPENDENT REASONABLE ASSURANCE REPORT – CLOSED-LOOP CONTAINMENT

We have performed a reasonable assurance engagement on the operating effectiveness of [Institution]'s Closed-Loop Containment architecture, comprising endpoint isolation, identity-tier-1 token revocation, and network-segment cutover, against the suitable criteria in Schedule A, for the year ended [date].

In our opinion, [Institution]'s Closed-Loop Containment operated effectively, in all material respects, in accordance with the criteria in Schedule A, throughout the year ended [date], including: (a) named decision rights signed by accountable executives; (b) override path tested under simulated false-positive on [n] occasions; (c) evidence-chain emission with cryptographic integrity for each closed-loop action; (d) compliance with the EU AI Act Article 14 human oversight requirements.

[Practitioner firm], [Partner name], [Date]
Engagement performed under ISAE 3000 (Revised).

Audit-grade close-out

Unattested provable autonomy is an internal claim. Attested provable autonomy is an institutional posture. The template above is the prerequisite for the underwriter pack (Paper 14), the supervisory engagement pack (Paper 13), and the M&A diligence pack (Paper 14 IV-bis). The audit cycle and capital cycle now coincide.

— SECTION —

X-novies. Operational artefacts — closed-loop safety case, misfire scenarios, and regulator cross-examination

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

Provable autonomy requires a safety case the institution can defend at supervisor challenge. The three artefacts below name the operational guarantees: per-action safety preconditions, named misfire scenarios with engineered responses, and a regulator cross-examination table the institution rehearses before it operates.

Action class	Pre-conditions	Blast-radius cap	Rollback guarantee	Kill-switch latency
Endpoint isolation	Confidence ≥ 0.85; quorum 2-of-3; not on no-isolate list	Single host; affects only host's network reachability	Auto-rollback ≤ 30 min if no Sev-1 confirm	≤ 60s manual override
Identity tier-1 token revocation	Confidence ≥ 0.90; quorum 2-of-3; user not on protected list	Single user; sessions invalidated; account not deleted	Auto-rollback ≤ 30 min if no Sev-1 confirm	≤ 60s manual override
Network segment cutover (contained host)	Host already isolated; segment not customer-facing	Single segment; pre-mapped failover available	Auto-rollback ≤ 30 min	≤ 60s manual override
IAM role suspension (privileged)	Human approval gate; not autonomous	Single role; sessions invalidated	Manual restore (no auto-rollback)	Immediate (gate is human)
AD account deletion / key destruction	Approval-gated; never autonomous	—	Non-reversible (CISO + GC sign-off)	—

Misfire scenario appendix — four scenarios the institution drills quarterly:

Scenario	Detection signal	Engineered response
False-positive containment of legitimate process	Auto-rollback fires after 30 min if no Sev-1 confirmation; analyst review surfaces FP via dashboard within 5 min	Service restored ≤ 35 min from action; FP recorded in register; rule re-tuned within 30 days
Failed rollback (e.g. orchestration tool unavailable)	Rollback monitor detects no acknowledgement within 90s	Pages on-call CISO + Detection Eng Lead; manual rollback procedure executes; escalation to COO if no manual rollback within 15 min
Conflicting telemetry (signals contradict)	Validator quorum not reached (e.g. 1 of 3 sources affirms; 2 disagree)	Closed-loop blocks; alerts to analyst queue with full evidence package; human authorisation required to proceed
Override path failure (override request not honoured)	Override request timeouts (>60s) detected and logged	Failover to secondary override channel; incident logged as Sev-1 control failure; post-incident review mandatory

Regulator cross-examination table — questions a supervisor can be expected to put to the institution, with the evidenced answer for each.

Supervisor question	Evidenced answer
---------------------	------------------

Who authorised this autonomous action?	Decision-Rights Register entry DR-CONT-EP-001, signed by CISO, CRO, COO, GC; version-controlled; reviewed quarterly
How was human oversight ensured (EU AI Act Art. 14)?	Override path tested in [drill records]; post-action human review within 24h; decision rights pre-signed by named human
What is the action's reversibility?	Auto-rollback within 30 min if no Sev-1 confirm; rollback drilled [n times]; mean rollback time [X] minutes; max [Y]
What evidence chain proves this action operated correctly?	Hash-chained signed JSON manifest; retained 5 years under DORA Art. 17; ISAE 3000 attested annually
What controls prevent customer-facing impact?	Action class matrix excludes customer-facing service shutdown from autonomous scope; blast-radius cap operates per-action; kill-switch latency < 60s

Operational artefact

Provable autonomy is now operationally defensible: the safety case names every guarantee, the misfire appendix names every drilled failure mode, and the regulator cross-examination is rehearsed before the supervisor asks. The institution that operates closed-loop without these artefacts is not yet defensible; with them, the supervisor engagement is a verification rather than a discovery.

— 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] Åström, K. J. & Murray, R. M. *Feedback Systems: An Introduction for Scientists and Engineers*, 2nd ed. Princeton, 2021.
- [02] European Commission. *Final assessment of the AI Act conformity for high-risk autonomous systems. EU AI Act guidance*, 2024.
- [03] OECD. *Recommendation of the Council on Artificial Intelligence*. OECD/LEGAL/0449, 2019 (rev. 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 05 — local data window

Paper 05 uses closed-loop deployment data from 9 institutions in the dataset that have deployed reversible-action closed-loop containment in production; time window Q1 2024 to Q1 2026. Settling-time data are summary statistics from each institution's drill records (no raw telemetry). Sample is small and skews toward Tier-1 financial services; generalisation to other sectors should be tested in the institution's own context.

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 The Autonomous SOC Is No Longer Optional — Doctrine Paper 05 of 20 • v6.1 — Operational Artefacts (Polished).