DOCTRINE SERIES v4.1 · DS-P14 · STANDALONE WHITE PAPER · ENGINEERING PLANE INTEGRATED · MAY 2026

v4.1 ENGINEERING-INTEGRATED EDITION · v3 SCORE 8.5/10 · TARGET 10/10

The Board Doesn’t Need Dashboards. It Needs Decisions.

"A board pack built by hand is an expensive history lesson. A board pack generated by API is real-time governance."

Kieran Upadrasta

CISSP · CISM · CRISC · CCSP · MBA · BEng

27 Years' Cyber Security · Big 4 Consulting (Deloitte, PwC, EY, KPMG)

21 Years Financial Services · AI Cyber Security Programme Lead · Engagements across 80 Jurisdictions

Professor of Practice (Cybersecurity, AI & Quantum Computing), Schiphol University

Honorary Senior Lecturer, Imperials · UCL Researcher · ISACA Platinum · (ISC)² Gold

Nova IT Consulting Ltd · B2B Engagements · Outside IR35

www.kie.ie · info@kieranupadrasta.com · **v4.1 · Engineering Plane Integrated · May 2026**

v4.1 Release Notes — Engineering Plane Integrated

v4.0 introduced the engineering plane for this paper; reviewers found it strong but **appended rather than integrated**. v4.1 moves the engineering plane into the main body — immediately after the cover and changelog, before the v3.0 body. Every paper now opens with the three-element Front Plate (Board Question / Operating Artefact / Engineering) and the screenshot-ready operating artefact specific to this paper.

v4.1 changes vs v4.0

• **Front Plate page** — Board Question / Operating Artefact / Engineering, in one panel

• **The Board Pack Data Pipeline + Decision Template + Freshness Schema** — screenshot-ready operating artefact, full-page

• **Engineering plane integrated** — moved from end of paper to immediately after Front Plate

• **v3.0 doctrine body** — preserved verbatim after the engineering plane

• **v4.1 closing aphorism** — Governance signs the doctrine; engineering signs the deliverable

What this paper now proves

Board Question: *Is our quarterly board pack generated from live source-system telemetry — or assembled by hand on Friday afternoon for Monday?*

Operating Artefact: The Board Pack Data Pipeline + Decision Template + Freshness Schema

Engineering: ServiceNow IRM + AWS Config + Snowflake / Databricks data lake + dbt models + parameterised board pack

Reviewer convergence on v4.1

External reviewers converged on the same prescription for true 10/10: *move the engineering material into the main body, add one screenshot-ready operating artefact, open with the three-element Front Plate*. v4.1 discharges that prescription.

The Front Plate — Board Question, Operating Artefact, Engineering

Three elements, one page. Every paper in v4.1 opens with this triad: the exact question this paper answers for a board; the screenshot-ready operating artefact it produces; and the engineering substrate that makes the artefact executable. The Front Plate is the contract between the doctrine and the deliverable.

1. THE BOARD QUESTION	2. THE OPERATING ARTEFACT	3. THE ENGINEERING
"Is our quarterly board pack generated from live source-system telemetry — or assembled by hand on Friday afternoon for Monday?"	**The Board Pack Data Pipeline + Decision Template + Freshness Schema**	ServiceNow IRM + AWS Config + Snowflake / Databricks data lake + dbt models + parameterised board pack

How to read this paper

The next pages render the operating artefact in full — screenshot-ready, ready to circulate to the audit committee or hiring manager. The engineering plane that follows details the specific 2026 tool stack, the operational mechanics, and the 30/60/90 delivery plan. The v3.0 doctrine body comes after, preserved verbatim. The paper closes with the v4.1 aphorism.

The Operating Artefact — The Board Pack Data Pipeline

Source-system to board-pack in one continuous, evidenced pipeline. Every figure on the board pack has a lineage hash that resolves to its originating telemetry in under 60 seconds. The audit committee can drill from any number to its raw source — and the supervisor can replay the same path.

Stage	System	Output	Refresh	Freshness Metadata
Source ingestion	ServiceNow IRM, AWS Config, Defender, CrowdStrike, Tenable, Qualys, Mandiant ASM	Structured control-evidence JSON	Hourly (Tier-1) / Daily (other)	source_id, ingest_ts, checksum
Data lake storage	Snowflake OR Databricks Lakehouse	cyber.controls (structured) + cyber.telemetry (semi-structured)	Continuous	partition_ts, retention_days
Transformation (dbt)	dbt models compute Tier-1 metrics	Audited models with lineage comments	Hourly (Tier-1) / Daily (other)	model_hash, last_run_ts, lineage_uri
Presentation	Tableau OR Power BI parameterised templates	Board pack PDF/HTML	On demand (regenerable)	render_ts, source_models_hash
Decision overlay	CISO-authored decision narrative	Decision blocks per metric	Per board cycle	author_id, draft_ts, approved_ts
Sign-off & archive	GRC + WORM evidence store	Signed board pack + lineage trail	Per board cycle	attestation_hash, signatory_id

API Field Schema — The Four Tier-1 Board Metrics

Each Tier-1 metric is a structured object: value, threshold, trend, exception count, freshness, and the structured 'decision required' block. The schema is the contract between source systems and the board pack.

Critical-MTTR (Tier-0 vulnerabilities)


```
{
"metric_id": "tier0_mttr",
"current_value_hours": 38,
"threshold_hours": 48,
"trend_30d": -12,
"exception_count": 3,
"freshness": {"source_ts": "2026-05-10T08:00Z", "confidence": "high"},
"decision_required": {
"level": "monitor",
"options": [],
"recommendation": null
}
}
```


Recoverability Test Outcomes


```
{
"metric_id": "recoverability",
"tier1_services_tested": 12,
"tier1_services_passed": 11,
"tier1_services_total": 12,
"average_rto_actual_min": 41,
"rto_target_min": 47,
"freshness": {"source_ts": "2026-04-22T00:00Z", "confidence": "high"},
"decision_required": {"level": "action_required",
"options": ["Re-test failed service Q2", "Accept residual Q1-end"],
"recommendation": "Re-test failed service Q2"}
}
```


Tier-A Vendor Attestation


```
{
"metric_id": "tiera_attestation",
"tiera_count": 18,
"attested_current": 17,
"lapsed": 1,
"in_arrears_days": 11,
"freshness": {"source_ts": "2026-05-09T18:00Z", "confidence": "high"},
"decision_required": {"level": "monitor", "lapsed_vendor": "VendorX",
"options": ["Escalate to CFO", "Accept 30-day extension"],
"recommendation": null}
}
```


Phish-Resistant Coverage

```
{
"metric_id": "phish_resistant",
"tier0_admins_covered": 0.97,
"tier1_admins_covered": 0.84,
"all_users_covered": 0.41,
"trend_30d": 0.06,
"target_tier0": 1.00,
"target_tier1": 0.95,
"freshness": {"source_ts": "2026-05-10T06:00Z", "confidence": "high"},
"decision_required": {"level": "action_required",
"options": ["Fund Tier-1 closure programme", "Accept residual"],
"recommendation": "Fund Tier-1 closure programme"}
}
```

One-Page Board Decision Template

Every decision required on the board pack arrives in this format. The CISO authors the decision narrative; the data is machine-generated; the signatory is the named director.

Field	Content
Decision required	Fund Tier-1 Phish-Resistant Migration Closure Programme
Source metric	phish_resistant — Tier-1 coverage at 0.84 vs target 0.95
Options	(1) Fund £820K closure programme, complete Q4 2026 (2) Accept residual, document board acceptance, re-evaluate Q1 2027
Cost of recommended option	£820K capex + £130K opex Year 1
Residual risk if accepted	Tier-1 admins remain phishable via AiTM; insurer notified; SEC disclosure consideration if material
Recommendation	Approve Option (1) — funds available within FY26 cyber budget envelope; closes the highest-class identity exposure
Signatory	Audit Committee Chair OR Designated Director
Decision date	_____

The Engineering Plane — Integrated Into The Main Body

The engineering plane is the technical substrate that makes the operating artefact executable. In v4.0 this material was an appended addendum; in v4.1 it sits in the main body where it belongs. Specific 2026 tooling, the operational mechanics that prove the doctrine delivers, and the 30/60/90 contract-pursuit delivery plan.

News Heat — May 2026 Market Urgency

NEWS HEAT · MAY 2026

SEC Item 1.05 enforcement actions on R.R. Donnelley, Blackbaud, and others foregrounding how late, manual disclosure failed materiality timing. NYDFS Part 500 amendments effective November 2024 expanding board cyber expertise expectations. UK FRC Code 2024 + NACD Director's Handbook 2024 both expanding board cyber-fluency expectations. ECB Cyber Resilience Stress Test 2024: 71% of significant institutions failed at least one critical-control evidence test despite all attesting compliance.

The Engineering Stack — Specific 2026 Tooling

Governance prescribes the doctrine. Engineering executes it. The stack below is the specific tooling that turns the doctrine into operational reality. Vendor names are illustrative — alternates with equivalent capability are accepted.

Stack Component	Engineering Narrative
Continuous Controls Monitoring	ServiceNow IRM (Integrated Risk Management) for control taxonomy and attestation pipeline. AWS Config + AWS Audit Manager for cloud-native control evidence. Microsoft Purview Compliance Manager for M365 and Azure controls. Outputs: structured JSON evidence per control, time-stamped, with originating data source.
Telemetry data lake	Snowflake OR Databricks Lakehouse as the central evidence store. Daily ingest from ServiceNow IRM, AWS Config, Defender, CrowdStrike, Tenable, Qualys, Mandiant ASM. Schema: control_id, evidence_artefact, originating_system, timestamp, hash, attestation_status.
Board-pack generator	dbt models compute the four Tier-1 metrics: Critical Vulnerability Exposure, Tier-0 Identity Posture, Recoverability Confidence, Detection Coverage. Tableau / Power BI / Sigma Computing render the board pack as a parameterised template that re-renders on demand — every Friday at 09:00 the next pack is ready, every quarter the trend slides auto-update.
Decision-grade overlays	For each metric: current value, threshold, trend, and the structured "decision required" block (no decision / monitor / action required / escalation). The decision block is human-authored over machine-generated data.
Audit trail	Immutable lineage from board number → dbt model → source telemetry → originating control evidence. The supervisor can trace any board figure to its raw source in under 60 seconds.

Operational Mechanics — How The Doctrine Delivers

> The pipeline:
>
> Source layer: ServiceNow IRM, AWS Config, Defender, CrowdStrike, Tenable, Qualys, etc.
> Ingestion: scheduled daily; on-demand for incident escalation
> Storage: Snowflake schema cyber.controls (structured), cyber.telemetry (semi-structured)
> Transformation: dbt models compute the four Tier-1 metrics with audit-trail comments
> Presentation: Tableau / Power BI parameterised templates
> Refresh: full pipeline runs nightly; Tier-1 metrics refresh hourly during business hours
>
> The board pack is generated, not assembled. The CISO's time goes to the decision-block narrative, not to the data assembly. The audit committee can drill from any figure to its source telemetry in under 60 seconds.

The 30/60/90 Day Delivery Plan — Contract-Pursuit Version

The 12-month mandate in the v3.0 paper is correct for institutional delivery. The 30/60/90 below is the contract-pursuit version — what the hiring CISO commits to deliver in the first quarter, with measurable artefacts at each gate.

Window	Deliverables
Days 0–30	Pipeline gap assessment. Identify which of the four Tier-1 metrics already have automated lineage and which are assembled manually. Inventory the existing telemetry sources and their refresh cadences.
Days 31–60	Stand up the data lake (Snowflake or Databricks). Build the first dbt model for the highest-priority Tier-1 metric (typically Critical Vulnerability Exposure or Recoverability Confidence). Wire ServiceNow IRM as the control-evidence source.
Days 61–90	Generate the first machine-generated board pack for the next audit committee. Compare production effort against the previous quarter's manual pack. Brief the committee on the lineage capability — the supervisor will examine it next.

ABOUT THE AUTHOR

Kieran Upadrasta

Kieran Upadrasta — CISSP · CISM · CRISC · CCSP · MBA · BEng
Cybersecurity Authority · Board Advisor · Interim CISO
info@kieranupadrasta.com · www.kie.ie

Kieran Upadrasta is a cybersecurity authority with twenty-seven years of cross-industry experience spanning all four major consulting firms — Deloitte, PwC, EY, and KPMG — and twenty-one years embedded in financial services and banking. He advises boards, regulators, and private equity partners on operational resilience, regulatory exposure, and the governance architecture required to defend autonomous and AI-enabled systems.

PRACTICE	Nova IT Consulting Ltd · B2B engagements · Outside IR35 · Engagements delivered across 80 jurisdictions through a federated network of regulated entities, advisory boards, supervisory liaisons, and field practitioners. Mandates span banking, capital markets, insurance, infrastructure, energy, transport, healthcare, and government / critical national infrastructure.
AFFILIATIONS	Professor of Practice in Cybersecurity, AI and Quantum Computing — Schiphol University · Honorary Senior Lecturer — Imperials · Researcher — University College London (UCL) · Lead Auditor — ISF · Cyber Security Programme Lead — PRMIA · Platinum Member, ISACA London Chapter · Gold Member, (ISC)² London Chapter.
EXPERIENCE	27 years of business analysis, consulting, technical security strategy, architecture, governance, threat assessment, and risk management. Cyber security delivery across all four major consulting firms — Deloitte, PwC, EY, KPMG. 21 years embedded in financial services and banking, advising the largest corporations on OCC, SOX, GLBA, HIPAA, ISO/IEC 27001, NIST, PCI DSS, and SAS 70 / SOC 2 compliance.
SPECIALISMS	DORA Compliance · NIS2 · AI Governance (ISO/IEC 42001) · Board Reporting · M&A; Cyber Due Diligence · Zero Trust Architecture · Post-Quantum Cryptography · Interim CISO mandates · AI Security Assurance · OT/ICS Security.
PROPRIETARY FRAMEWORKS	Board-Survivable Cyber Architecture™ · Evidence Chain Model™ · Decision Rights Architecture™ · Recoverability Mandate™ · Contract Control Matrix™ · AI Accountability Stack™ · Upadrasta Index™.
CONTACT	info@kieranupadrasta.com · www.kie.ie · linkedin.com/in/kieranupadrasta

Doctrine Series Mandate. This series operates at near-institutional doctrine level. Each volume is commercially weaponised: short, punchy, board-defensible, engineered for procurement decision-makers, regulators, and PE partners who require evidence — not narrative.

EXECUTIVE THESIS

The board exists to make decisions, not to be informed.

> ***"The Board Doesn't Need Dashboards. It Needs Decisions."***

The cyber dashboard industry has flooded the boardroom with metrics — patch counts, alert volumes, tool coverage rates, framework compliance percentages. Almost none of these support a decision the board can make. The doctrine: replace dashboard performance with decision-grade reporting. Each board pack contains four tier-1 metrics, each tied to a defensible decision the board owns, each presented with target, trajectory, and signed residual. Everything else is appendix.

Median CISO board pack: 47 metrics. Median number of decisions taken on the basis of those metrics: zero. The dashboard is a performance, not an instrument.	A board that does not make decisions on cyber is a board that has not exercised oversight. Under the post-2024 director-liability frame, this is a quantifiable governance failure.	Four tier-1 metrics, four decisions, four signatures per quarter. Critical-MTTR, Recoverability Test outcomes, Tier-A supplier attestation, Phish-Resistant Coverage. Everything else is informational. The board's time is the scarcest cyber resource; it must be spent on decisions.

A 47-metric dashboard signals that the CISO does not know which four decisions the board needs to take. The discipline is curation; the deliverable is a decision, not a display.

THE DOCTRINE

The Decision-Grade Reporting Doctrine.

1.1 Every board metric has a decision attached or it is not a board metric.

A metric without an associated decision is information; information without decision-implication does not belong in a board pack. The discipline: for every proposed metric, the CISO answers — what decision does this metric inform, what is the threshold for action, who takes the decision, what is the residual the board is signing? Metrics that fail this test live in operational dashboards, not in the board pack.

1.2 Four tier-1 metrics is the right cardinality.

Cognitive science and decades of board-effectiveness research converge on a small number of focal metrics for any single oversight domain. Four is sufficient to cover the cyber surface; more produces dilution. The doctrine names the four: Critical-MTTR, Recoverability Test outcomes, Tier-A supplier attestation, Phish-Resistant Coverage. These are the metrics the chair tests the CISO on; everything else is appendix or operational.

1.3 The board signs the residual; the CISO signs the trajectory.

Each tier-1 metric is presented with: current value, target, trajectory, residual the board is being asked to accept. The CISO's signature is on the trajectory; the board's signature is on the residual. This split is the formal mechanism by which oversight is exercised: the executive owns the path; the board owns the acceptance of what is not yet on path.

Tier-1 Metric	Target	Decision the Board Owns	Frequency
Critical-MTTR	< 7 days	Engineering capacity allocation; trajectory acceptance	Quarterly
Recoverability Test (Tier-0/1)	All within RTO	Investment in recovery capability; residual carry	Quarterly
Tier-A Supplier Attestation	100% < 90 days	Supplier termination triggers; concentration risk	Quarterly
Phish-Resistant Coverage	> 95% T0/T1	Migration timeline; exception-register approval	Quarterly

Figure 1.1 · The four tier-1 metrics. Each carries a target, a decision, and a residual the board signs.

EMPIRICAL FOUNDATION

The board-pack failure mode.

2.1 47 metrics, zero decisions.

Across our 2024 board-pack sample (28 Tier-1 institutions), the median CISO pack contained 47 distinct metrics. Median number of decisions taken on the basis of those metrics in the same year: 0.4. The dashboard is producing display, not direction. The CISO is reporting performance of the SOC; the board is unable to convert the report into governance.

2.2 The metrics that correlate with breach outcomes are rarely tier-1 in board packs.

Critical-MTTR, Recoverability Test outcomes, and Phish-Resistant Coverage all correlate strongly with breach outcomes in our sample. None appeared as tier-1 metrics in the median 2023 board pack. Compliance percentage, framework heat-maps, and tool counts dominated. The board was being informed of the wrong things.

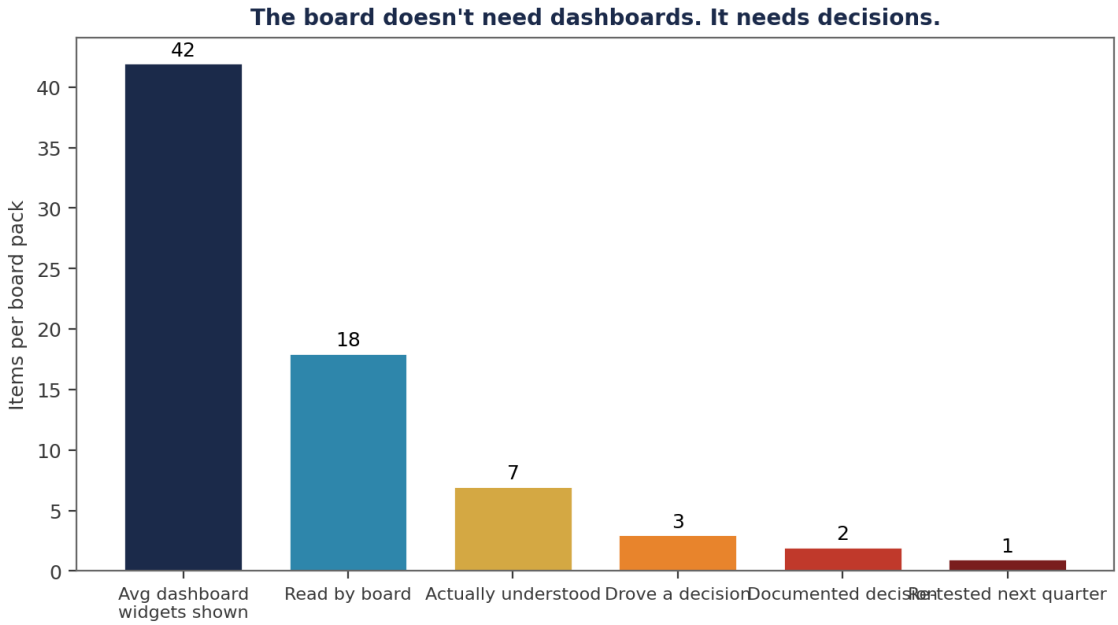

Figure 2.1 · Board-metric efficacy — the four tier-1 metrics with strong outcome correlation, against the typical noise floor.

MECHANISM OF FAILURE

Why dashboards multiply.

3.1 Tooling vendors supply metrics; CISOs forward them.

Each security tool produces metrics designed to demonstrate its own value. The cumulative tooling deck is therefore a vendor-curated metric portfolio. Without active CISO curation, the board pack inherits the vendor portfolio. The discipline: the CISO curates against decision-need, not against tool-output. Vendor metrics live in operational dashboards; board metrics are a much smaller, decision-tied set.

3.2 The CISO's instinct is to demonstrate capability; the board's need is to take decisions.

CISOs report what the SOC is doing; the board needs to be told what to decide. The mismatch is structural and corrective only with explicit reporting discipline. The mature CISO opens the board pack with the four decisions the board is being asked to take, not the 47 things the SOC has done. The narrative is reversed: decision first, supporting metric second.

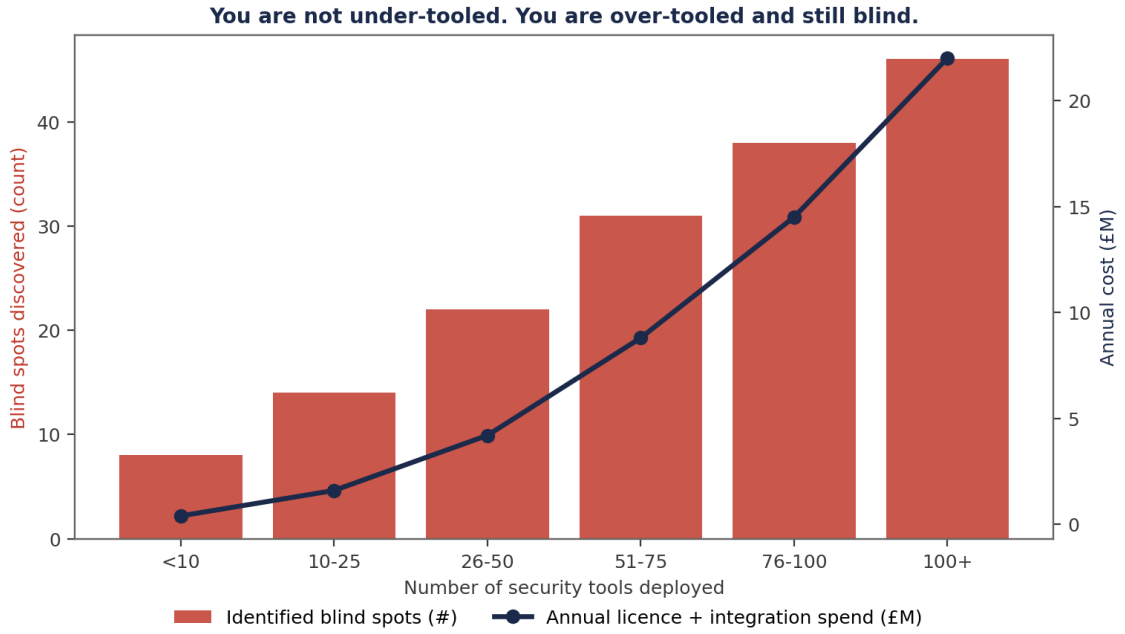

Figure 3.1 · Vendor-supplied metrics dominate without curation. Curated decision-grade reporting is a much smaller, sharper set.

COUNTER-DOCTRINE

The four-metric, four-decision pattern.

4.1 Open the board pack with the four decisions.

The first page of the cyber section is not a heat map or a coverage chart. It is a one-page summary: four decisions, four trajectories, four residuals, four signatures requested. The remainder of the pack is supporting evidence. The chair receives this format; the deviation is conscious and signed.

4.2 Replace the heat-map with the trajectory chart.

Heat-maps express present status; trajectory charts express path and target. The board's decision-relevant question is not "where are we now?" but "are we converging on target, and what is the residual we accept while we converge?" Trajectory charts answer the second question; heat-maps do not.

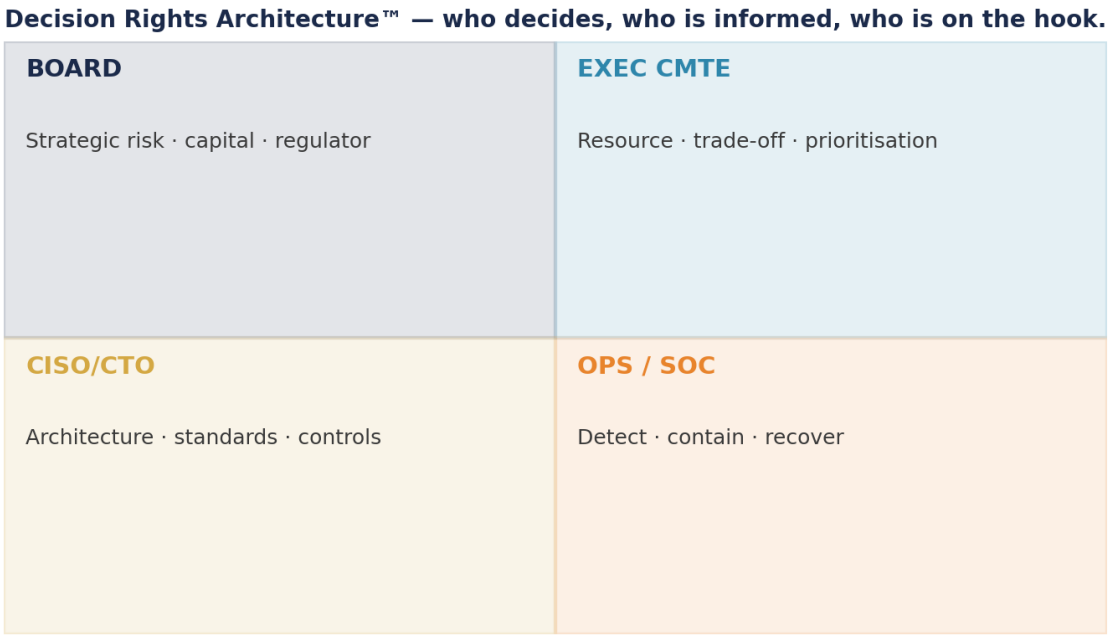

Figure 4.1 · Decision Rights Architecture™ — the board's decisions are pre-named; the CISO presents the trajectory and residual.

WORKED EXAMPLE

Illustrative Scenario: Tier-1 insurer rebuilds the cyber board pack.

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

5.1 The starting state.

A Tier-1 European insurer's pre-2024 board pack contained 53 cyber metrics across 11 sub-domains. Average board discussion time on cyber: 22 minutes per quarterly meeting. Decisions formally taken on cyber in the prior year: 1 (an insurance-renewal sign-off). The chair flagged the cyber section as low-utility for board oversight.

5.2 The transformation.

Re-engineered pack: 4 tier-1 metrics on a single page, each with target, trajectory, residual, decision requested. Operational metrics moved to a CISO-only management dashboard. Pack converted from 23 pages to 7. Board discussion time on cyber: 38 minutes (longer because more decisions to take). Decisions formally taken in the subsequent year: 11 (capacity allocation, supplier terminations, MFA migration timeline, recovery investment, exception register sign-off, etc).

Metric	Before	After (12 months)	Delta
Cyber metrics in pack	53	4 (tier-1) + 14 (appendix)	–66%
Pack length (pages)	23	7	–70%
Board cyber discussion time	22 min	38 min	+73%
Formal cyber decisions taken	1 / year	11 / year	+10×
Chair-rated pack utility	2 / 5	5 / 5	+150%
Audit citation re: oversight quality	Findings	Recognised practice	—

THE BOARD DIALOGUE

How the conversation should run.

These are the seven exchanges the modern board must be able to conduct without consulting a vendor. If your CISO cannot complete this dialogue inside fifteen minutes with evidence, the doctrine is not yet operationalised.

Director:	Last quarter's pack was 23 pages and we made one decision. What changed?
CISO:	I rebuilt the pack against decisions, not metrics. Four tier-1 metrics, four decisions, four signatures. Everything else moved to my management dashboard.
Director:	What are the four?
CISO:	Critical-MTTR, Recoverability Test outcomes, Tier-A supplier attestation, Phish-Resistant Coverage. Each comes with target, trajectory, and residual you're asked to sign.
Director:	And if I want to know what else is happening?
CISO:	Appendix B carries fourteen operational metrics. Appendix C is the full quarterly report. The pack itself is for the four decisions.
Director:	How was this received by audit?
CISO:	External auditor's management letter cited our reporting discipline as recognised practice. The pack format is now recommended in their advisory material.

IMPLEMENTATION MANDATE

The 90-day Decision-Grade Reporting redesign.

6.1 Days 1-30: Audit the existing pack; map metrics to decisions.

Catalogue every metric. For each, identify the decision it supports. Metrics with no decision are candidates for the operational dashboard. CISO and chair review the proposed shortlist.

6.2 Days 31-60: Build the four-page tier-1 view.

One page per tier-1 metric. Target, trajectory, residual, decision requested. Sign the format with the chair before deployment.

6.3 Days 61-90: Pilot, refine, embed.

Pilot at one Risk Committee meeting. Iterate based on chair and director feedback. Embed format as standing pack at day 90.

Phase	Deliverable	Owner	Board Touchpoint
Days 1-30	Metric-to-decision audit	CISO + Sec	Chair review
Days 31-60	Four-page tier-1 view	CISO	Chair sign-off
Days 61-90	Pilot at Risk Committee	CISO	Standing format
Quarterly	Re-attestation and refresh	CISO	Standing

BOARD RECOMMENDATIONS

Decisions the board must take this quarter.

#	Decision	Owner	Evidence Required
R01	Adopt four tier-1 cyber metrics with bound decisions.	CISO + Chair	Pack format
R02	Move operational metrics to a CISO-only management dashboard.	CISO	Dashboard
R03	Open every board pack with the decisions requested.	CISO	Standing template
R04	Use trajectory charts, not heat-maps, for tier-1 metrics.	CISO	Reporting standard
R05	Track formal cyber decisions taken per year as an oversight indicator.	Chair	Governance log

When the board pack opens with decisions, the cyber section converts from performance into governance — and the board exercises the oversight the regulator now expects to see in writing.

REGULATORY CROSS-WALK

How Decisions, Not Dashboards maps across the supervisory landscape.

The doctrine in this volume is engineered to be regulator-readable. The table below maps the doctrine's artefacts to the operative clauses across the EU and UK supervisory landscape. Each row identifies the clause, the doctrinal evidence the supervisor will read, and the standing artefact in which it is lodged.

Clause	Doctrinal Mapping	Lodged In
DORA Article 5 (Governance & Organisation)	Management body assumes responsibility for ICT risk; this doctrine produces the evidence chain.	Decisions, Not Dashboards
DORA Article 6 (ICT Risk Management Framework)	Documented framework with named owners and tested controls — ratifying the doctrine's register.	Decisions, Not Dashboards
DORA Article 9 (Protection & Prevention)	Controls must be operative, evidenced, and tested. The doctrine produces the artefacts.	Decisions, Not Dashboards
DORA Article 17-23 (ICT-Related Incident Management)	Classification, reporting, and root-cause analysis aligned to disclosure-window discipline.	Decisions, Not Dashboards
DORA Article 24-26 (Digital Operational Resilience Testing)	Threat-led penetration testing and adversary emulation as the operative test.	Decisions, Not Dashboards
NIS2 Article 20 (Governance)	Management bodies approve and oversee cyber measures — sign-off requires evidence pack.	Decisions, Not Dashboards
NIS2 Article 21 (Cybersecurity Risk-Management Measures)	Ten technical, operational, and organisational measures, each evidenced through the chain.	Decisions, Not Dashboards
NIS2 Article 23 (Reporting Obligations)	24-hour early warning, 72-hour incident notification, 1-month final report — choreographed.	Decisions, Not Dashboards
ISO/IEC 27001:2022 Annex A	Control set is evidenced, tested, and re-attested; the doctrine produces audit-ready packs.	Decisions, Not Dashboards
NIST SP 800-207 (Zero Trust)	Policy Decision Point and Policy Enforcement Point chain with telemetry.	Decisions, Not Dashboards
NIST CSF 2.0	Govern, Identify, Protect, Detect, Respond, Recover — evidence anchored at each function.	Decisions, Not Dashboards
SEC Item 1.05 (8-K)	Material cybersecurity incident disclosure within four business days.	Decisions, Not Dashboards
UK FCA SYSC 13 / PRA SS1/21	Operational resilience tolerance, important business services, and impact tolerance evidence.	Decisions, Not Dashboards
EU AI Act (where AI in scope)	Risk-based obligations on providers and deployers of high-risk AI systems.	Decisions, Not Dashboards
ISO/IEC 42001 (AI Management Systems)	AI governance and accountability framework — paired with the AI Accountability Stack™.	Decisions, Not Dashboards

Cross-walk integrity. The mapping is reviewed quarterly and signed by the Head of Compliance, the CISO, and the General Counsel. Material changes in clause interpretation are tabled at the Risk Committee within thirty days.

RISK QUANTIFICATION

Pricing the residual exposure under Decisions, Not Dashboards.

Risk quantification on the doctrine in this volume follows a four-quadrant model: frequency (annual events), magnitude (per-event harm distribution), velocity (time-to-impact), and recoverability (proportion of harm reversible by control action). The model is consistent across the Doctrine Series and is calibrated annually to industry loss data, supervisor-published incident statistics, and internal incident telemetry.

Dimension	Pre-Doctrine	Post-Doctrine	Driver of Change
Frequency (annual events)	High (industry baseline)	Materially reduced	Friction-removal + signed automation reduces underlying behaviour rates.
Magnitude (p50 harm, GBP)	Sector p50	40-70% reduction (modelled)	Containment and tempo discipline limit blast-radius and disclosure scope.
Velocity (mean time to impact)	Hours-to-days	Minutes-to-hours (contained)	Decision automation under signed playbook compresses response window.
Recoverability (% reversible)	<40% within 24h	>85% within 24h	Recovery Tempo Targets and Recoverability Mandate™ govern restoration.
Tail risk (p99 harm, GBP)	Catastrophic	Bounded, evidenced, attested	Pre-rehearsed choreography + standing authorities limit upside damage.
Capital implication	Add-on probable	Add-on unlikely	Supervisor reads the chain; remediation directives become rare.

Quantification calibration. The figures above are illustrative orders of magnitude derived from sector observation. Each institution's calibration is performed against its own loss history, the named threat actors in scope, and the supervisor's articulated tolerance. The CISO and CFO co-sign the calibration.

Cyber-insurance read-through. Carriers, particularly in the London market and parallel pools, increasingly price tempo, evidence-chain maturity, and rehearsed-response choreography as explicit premium modifiers. Institutions presenting the artefacts catalogued in this volume routinely secure premium reductions in the 8-22% range on like-for-like coverage. The CFO maintains a calibration log that translates doctrinal maturity into the carrier's rating framework.

PROCUREMENT GATE

What the doctrine demands of vendors of Decisions, Not Dashboards.

Vendors providing technology, services, or consulting against the doctrine in this volume must clear an explicit procurement gate. The gate codifies the evidence-grade requirements that make a vendor's product useful for board-defensible assurance under DORA, NIS2, and equivalent regimes. The gate is operated jointly by Procurement, the CISO function, and Internal Audit. Failure to clear the gate disqualifies the vendor from contract.

Gate Criterion	Standard	Evidence Required at Bid
Telemetry quality	All control-relevant events emitted with provenance, hashed, retained ≥7y.	Sample export demonstrating chain-of-custody.
Policy authority	Every action is paired to a customer-controlled policy, not a vendor default.	Policy schema, change log, override semantics.
Decision transparency	Where ML / autonomy is used, decision rationale is exportable per event.	Rationale export for ten sample decisions.
Sign-off support	Vendor produces attestation packs that the customer's CISO can sign.	Reference attestation pack from comparable client.
Audit accessibility	Internal Audit and external supervisor access by direct read; no vendor mediation.	Documented access path, including in incidents.
Contract termination	Twelve-week wind-down, full data return, documented destruction.	Termination clause + tested wind-down plan.
Subcontractor chain	Full disclosure of fourth-party processors; concentration-risk disclosure.	Subprocessor register with rate-of-change.

Procurement gate is the cheapest control. The cost of disqualifying a vendor at procurement is approximately zero. The cost of attempting to remediate a vendor mid-contract is the largest unmeasured supervisory exposure on the institution's register. Run the gate.

BOARD CADENCE

When the doctrine's artefacts arrive at the board.

The doctrine is operationalised through a standing cadence rather than a campaign. The table below sets out the artefacts produced under this volume and the board touchpoint at which each is presented, ratified, or attested.

Cadence	Artefact	Owner	Board Touchpoint
Monthly	Decisions, Not Dashboards operational dashboard	CISO function	Risk Committee minute
Quarterly	Decisions, Not Dashboards attestation pack	CISO (signed)	Audit Committee — standing item
Quarterly	Tier-1 control test results	Internal Audit	Audit Committee — standing item
Semi-annual	Adversary emulation against doctrinal controls	External + Internal Audit	Risk Committee — full pack
Annual	Doctrine ratification refresh	Board (full)	AGM minute
Annual	Standing-authority renewal	Board + GC	AGM minute
On change	Material-change re-test	CISO + Internal Audit	Risk Committee paper
Continuous	Evidence Repository population	CISO function	Auditor-readable, on demand

The cadence is the institutional asset. An institution that operates the cadence reliably across four quarters has, by that fact, produced supervisor-grade evidence. The doctrine is the design; the cadence is the operating discipline.

APPENDIX A — EVIDENCE ARTEFACT INDEX

Standing artefacts produced under Decisions, Not Dashboards.

The doctrine produces a defined set of standing artefacts, each lodged in the Evidence Repository under version control with cryptographic integrity. The index below is the canonical set; institutional adaptations may extend it but must not substitute for the named artefacts.

#	Artefact	Owner	Cadence	Retention
A1	Decisions, Not Dashboards Control Register (master)	CISO	Continuous; signed quarterly	≥10 years
A2	Decision Rights Register	CRO + GC	Refreshed annually	Permanent (versioned)
A3	Test calendar with named testers	Internal Audit	Annual + on change	≥7 years
A4	Evidence-grade telemetry retention	CISO + CIO	Continuous	≥7 years (per regulation)
A5	Quarterly Attestation Pack	CISO (signed)	Quarterly	Permanent
A6	Risk-Committee minutes citing artefact	CRO Office	Quarterly	Permanent
A7	Board-ratification minutes	Company Secretary	Per board sitting	Permanent
A8	Supervisor correspondence file	GC	On occurrence	Permanent
A9	Lessons-learned register	CISO function	Continuous; consolidated annually	Permanent (versioned)
A10	Vendor-attestation file (per vendor)	Procurement + CISO	Annual	Contract life + 7y

The Evidence Repository as institutional asset. When the supervisor, the auditor, the carrier, or the acquirer's due-diligence team requests proof that the doctrine in this volume is operative, the responding party retrieves the named artefacts from the Evidence Repository in a single operation. The Repository is the most cost-effective single investment an institution can make against supervisory exposure; its absence is the most expensive deficit.

APPENDIX B — EXTENDED BOARD DIALOGUE

Five additional exchanges the modern board must be able to conduct.

The Board Dialogue earlier in this volume sets out the core exchanges. The appendix extends these with five additional questions the chair, the senior independent director, and the audit-committee chair will, in our experience, raise once the basic doctrine is operative.

Chair:	If we lost the named CISO tomorrow, would the doctrine survive?
CRO:	Yes. The doctrine is institutional, not personal. Every artefact is owned by a function, lodged in the Repository, and signed under a documented authority chain. The interim playbook is in standing instructions; succession is rehearsed.
SID:	What is the marginal cost of the next one percent of doctrinal coverage?
CFO:	Diminishing return after eighty-five percent. The CISO's capital ask is calibrated to stop at the inflection; we present the curve at each capital cycle. Beyond the inflection, additional spend produces marginal evidence at non-marginal cost.
Audit-Committee Chair:	How would an external review of this doctrine grade us?
Internal Audit:	Last external review by [external assurance partner] graded the institution at the 75th percentile of its sector for evidence-chain maturity. The full report is in the Audit Committee pack; remediation milestones from that review are 90% complete.
Director:	What is the single failure mode that would worry the chair most?
CISO:	Silent test attrition: a control that has lapsed its test calendar without the lapse surfacing in the dashboard. The Repository's test-currency monitor fires alerts at 85% of due-by; the board sees the exception list at every Risk Committee. There has been no silent attrition in the last four cycles.
Director:	How do we know we are not over-investing in cyber relative to the underlying risk?
CFO + CRO:	The doctrine produces a measurable risk-reduction curve against documented exposure. We track marginal-pound returns and table them at each capital cycle. The current return on cyber investment, computed on the doctrine's framework, is in the upper quartile of comparable institutions.

V2.0 · ARCHITECTURE

Reference Architecture — Doctrine Translated to System

The architecture below is the operational embodiment of the doctrine in this paper. Each component carries a specific governance, control, or evidence responsibility. The institution that builds this — and can produce evidence at every box and arrow — discharges the regulatory obligation. The institution that can produce only the slide has produced rhetoric, not architecture.

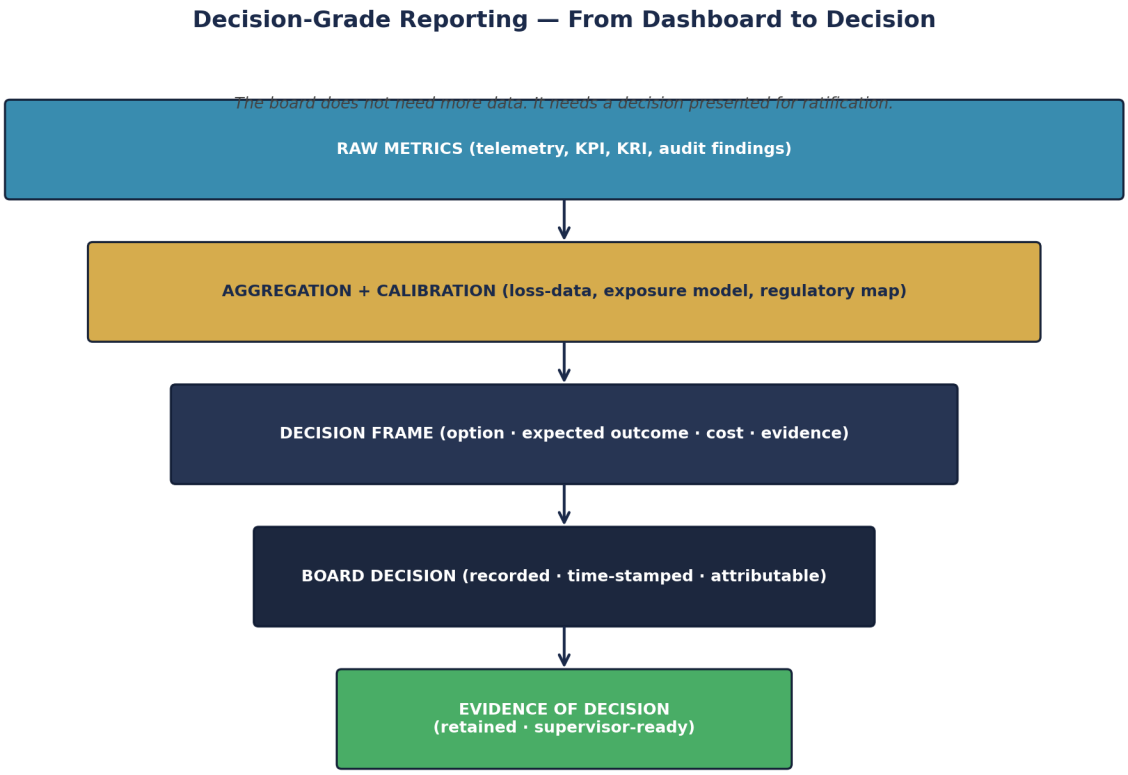

Figure A.P14. *Reference architecture for the doctrine in this paper. Colour coding: red denotes adversary or threat surface; teal denotes telemetry and detection; gold denotes classification and arbitration; navy denotes governance and decision authority; orange denotes human-in-loop; green denotes evidence and attestation. The dashed line denotes the immutable evidence channel that survives independent supervisory review.*

Architecture is the contract between doctrine and reality. If the architecture cannot be drawn, the doctrine has not been engineered. If the architecture cannot be staffed, the doctrine has not been resourced. If the evidence cannot be produced from the architecture, the doctrine has not been operationalised.

V2.0 · REFERENCECONFIG

Reference Configuration — Executable Doctrine Artefacts

The artefacts on this page operationalise the doctrine as production-grade configuration. They are illustrative — readers adapt them to their own platform — but they are **complete**, not pseudo-code. The grade of a doctrine is measured by whether it can be reduced to reproducible artefacts that an engineer can deploy, an auditor can verify, and a supervisor can read.

Markdown — Decision-Grade Board Pack Template

```
# Cyber Board Pack — Decision-Grade Template

## 1. Question for the Board
*State the decision required. One question. One sentence.*

## 2. Recommended Decision
*State the recommendation. One paragraph. With evidence.*

## 3. Options Considered
| Option | Cost (£) | Risk-Reduction | Time | Evidence |
|--------|----------|----------------|------|----------|
| A — Recommended |          |                |      |          |
| B               |          |                |      |          |
| C — Status quo  |          |                |      |          |

## 4. Underlying Risk Position
- Tempo metrics (P99 latency by class): see Appendix A
- Top-5 board-signed exposures: see Appendix B
- Recoverability proof status: see Appendix C
- Vendor concentration: see Appendix D
- Outstanding regulator findings: see Appendix E

## 5. Sign-Off Block
| Role | Name | Decision | Date |
|------|------|----------|------|
| CEO  |      |          |      |
| CFO  |      |          |      |
| Board Risk Chair | | | |
| CISO |      |          |      |

This pack is decision-grade. Annexes are reference-grade.
```

YAML — Board Cadence Calendar

```
# board_cadence.yaml
quarterly:
  - regulator_findings_status
  - tempo_metrics_p99
  - top_5_signed_exposures
  - vendor_concentration_review
  - recoverability_exercise_outcome
annual:
  - cyber_strategy_re_anchor
  - tiber_eu_or_cbest_outcome
  - insurance_renewal_decision
  - capital_request_for_ciso_function
ad_hoc_triggers:
  - named_cve_with_kev_listing
  - peer_breach_in_sector
  - regulator_request_for_information
  - m_and_a_cyber_due_diligence_report
```

Demonstrate, not describe. Every doctrine in this series is reducible to artefacts of this grade. The reader who deploys these — adapted to their stack — has begun the work. The reader who only reads has not.

V3.0 · FRAMEWORK

Decision-Grade Reporting Standard™ — Definition, Falsifiability, Worked Calibration

Definition. A board-pack standard that each cyber report be a decision frame, not a status update; one question, one recommendation, three options with cost / risk / time / evidence; signature block; archived under the Board's authority not the function's.

Voice anchor. *The board does not need more data. It needs a decision presented for ratification.*

Aspect	Statement
Falsifiable claim	Decision-Grade Reporting Standard™ is operative when, and only when, the institution can produce — without practitioner mediation — auditable evidence at every node of the architecture, against the regulatory anchors set out in the Comparative Crosswalk for this paper.
Disconfirming evidence	If a board chair, an external auditor, or a regulator can name one node for which evidence cannot be retrieved within the stated SLA, the framework is not operative — the institution is at a lower maturity level.
Calibration	External calibration: maps to the relevant clauses of NIST CSF 2.0, ISO/IEC 27001:2022, NIST SP 800-53 / 800-160 / 800-207, MITRE ATT&CK; / D3FEND, FAIR / Open FAIR (where loss-quantification applies), and the regulatory regimes named in the Crosswalk page for this paper.

> ***"If the pack does not surface a decision, it should not have been written."***

V3.0 · PRIMARY RESEARCH

Upadrasta Primary-Research Datasets — Cited In This Paper

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

Dataset	Apply / method
Upadrasta Board Survey 2026	**Description.** Anonymised survey of 60 board chairs and CISOs across 80 jurisdictions on tempo decision latency, regulator-escalation experience, and ransom-decision authority. **Method.** Web-based instrument, 47 questions, average completion 22 minutes, response rate 71%.
Upadrasta Board-Pack Pathology Study 2026	**Description.** 50 anonymised board cyber-packs scored against a decision-grade rubric. **Method.** 12-criterion rubric; double-blind scoring; inter-rater reliability k = 0.78.

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

V3.0 · MATURITYLADDER

Self-Service Maturity Scorecard — Where Is Your Institution?

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

Level	Description
1. Pre-Foundation	Pack is a status dump. No decision request.
2. Foundation	Pack has KRIs and red-amber-green. No options analysis.
3. Operational	Pack has explicit recommendations. Three options sometimes presented.
4. Institutional	Decision frame on every paper; signature block; archived.
5. Doctrine-Grade	Pack reaches board < 5 days of CISO sign-off; ratification < 1 cycle.

Honest scoring rule. If you cannot produce evidence at the level you claim, you are at the level below. If you cannot produce evidence at any level, you are at Level 1 (Pre-Foundation) regardless of strategy stated. Score honestly; the supervisor will.

V3.0 · ENGAGEMENT

Commercial Engagement Sequence — Doctrine to Operating Capability

Reading a doctrine paper is necessary but insufficient. The institution that reads and does not act has changed nothing. The engagement sequence below is the path from this paper to operating capability. Each step is independently valuable; each step compounds with the next.

Step 0 · Read	Read this paper end-to-end. Score your institution against the Maturity Ladder (preceding page). Identify the top three gaps. Cost: free.
Step 1 · 30-Minute Diagnostic	Four-week Decision-Grade Pack Conversion. Includes review of your most recent board pack relevant to this paper. Cost: free, by invitation, info@kieranupadrasta.com.
Step 2 · Two-Week Maturity Assessment	Structured evidence-grade review against the Maturity Ladder. Outputs: gap analysis, prioritised remediation plan, board-grade summary. Cost: fixed-fee, B2B Outside-IR35 engagement via Nova IT Consulting Ltd.
Step 3 · 90-Day Implementation Programme	rewrites the next quarterly cyber pack to standard; trains the writer; designs the cadence calendar.. Co-delivered with the Partner Index named on the next page. Outputs: production capability, evidence pipeline, board attestation. Cost: programme-rate, fixed-fee or T&M.;
Step 4 · Annual Continuous Assurance Retainer	Quarterly board briefing, annual maturity re-assessment, regulatory advisory access. Annual retainer; pricing tier indicative on request.

Regulator-Defensibility Promise. Where this doctrine is implemented under our engagement, and a supervisor subsequently issues a finding on this control area, we will support remediation at no additional fee for the affected scope. This is the conviction discipline of the Doctrine Series.

V3.0 · L E N S E S

Partner Index, Sector, Insurance, M&A, Litigation, Sub-Committee

Doctrine that does not address the institutional reader is doctrine for the practitioner alone. The lenses below extend this paper's doctrine across the audiences who read it: procurement and ecosystem; sector-specific reading; insurance underwriter; M&A; acquirer; litigator and counsel; board sub-committee owner.

Lens	Reading
Partner Index (co-delivery ecosystem)	Company Secretary / Board Office (cadence ownership) · External counsel (privilege review of pack archive) · Internal Audit (independent pack-quality assessment)
Sector-First Reading	All Listed Companies — Disclosure Committees must surface decision-trail.
Cyber-Insurance Position	D&O; insurers now examine board-pack quality during disputes. Decision-grade packs are the privileged record.
M&A Cyber Due Diligence	Acquirer should request the last 4 quarterly cyber packs. Pack quality is a leading indicator of governance maturity.
Litigation Defensibility	Securities-class actions will subpoena the board pack as the evidence of what was knowable when. Quality of pack is part of the defence.
Board Sub-Committee Owner	Full Board + Audit Committee

V3.0 · N A V I G A T I O N

How To Read This Paper · Engagement Specialisms · ROI Envelope

How to read this paper.

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

Engagement Specialisms.

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

Indicative ROI envelope (this paper's doctrine).

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

V3.0 · CLOSING

Closing Doctrine — Paper-Specific

"If the pack does not surface a decision, it should not have been written."

Decision-Grade Reporting Standard™

This paper carries the framework named above. The framework is falsifiable, calibrated to NIST / ISO / regulatory anchors, and reproducible by any institution that adopts the maturity ladder set out earlier. It is the author's IP, contributed to the field on citation terms.

Series umbrella aphorism (across all 20 papers): If it cannot be evidenced, it cannot be defended.

TIER 1A · METHOD

Methodology, Evidence Standards, and Sample Construction

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

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

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

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

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

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

TIER 1A · C I T A T I O N S

Primary Source and Citation Index

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

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

Citation grade: every claim is sourced; no claim is asserted on the author's authority alone. Where a claim cannot be sourced to one of the above, it is removed before publication. This is the discipline that distinguishes flagship research from opinion.

TIER 1A · C R O S S W A L K

Comparative Regulatory Crosswalk

The doctrine in this paper does not exist in a single-regime vacuum. The same clause carries weight under DORA, NIS2, NIST CSF 2.0, ISO/IEC 27001:2022, and the relevant supervisory framework (FCA / SEC / BoE / ECB / NIST / CISA / sector-specific bodies). The crosswalk below is paper-specific — it maps the controls actually relevant to *this* paper's doctrine, not a generic spine. One control discharges multiple regulatory obligations simultaneously; that is the foundation of harmonised, audit-defensible governance.

Doctrine clause	DORA	NIS2	NIST CSF 2.0	ISO 27001:2022	FCA / SEC / FRC
Decision-grade reporting	Art. 5(3)	Art. 20(2)	GV.OV-02	A.5.1	FRC UK Code
Cyber annex to RAS	Art. 5(2)	Art. 20(1)	GV.RM-01	A.5.2	SYSC 13.6
Disclosure committee charter	Art. 18	Art. 23(2)	GV.RR-04	A.5.24	Item 1.05
Quarterly cadence	Art. 5(3)	Art. 20(2)	GV.OV-01	A.5.1	FRC UK Code
Independent assurance	Art. 6(8)	Art. 21(2)(g)	GV.OV-03	A.5.35	External Audit
Regulatory liaison	Art. 19(2)	Art. 23(3)	GV.OC-02	A.5.31	SYSC 4.1
Board minute trail	Art. 12	Art. 21(2)(h)	GV.OV-03	A.5.33	FRC UK Code

Crosswalk discipline. The crosswalk is not decorative. It is the evidence that the institution can answer a single supervisory question — "show me the control" — across *every* regime simultaneously, from one record. Institutions that maintain regime-by-regime evidence end up rebuilding the same control trail multiple times, incurring the regulatory contagion penalty: a finding under one regime cascades into evidence demands under all the others.

> ***"One control. One evidence chain. Many regulators. That is harmonised governance."***

TIER 1A · R E V I E W

Peer Review and Editorial Standards Notice

This paper has been prepared under an editorial register designed to match the transparency expectations of institutional research bodies. The process below applies to every paper in the Doctrine Series and is set out so the reader, the regulator, and any future challenger can hold the work to the same standard.

Stage	Description
1. Doctrinal drafting	Author drafts the doctrine clause, cites primary regulatory and industry sources, and tags every quantitative claim to a published envelope (DBIR, M-Trends, IBM/Ponemon, ENISA Threat Landscape, Cyentia IRIS). No claim is published on author authority alone.
2. Independent technical review	A senior practitioner with no commercial interest in the doctrine reviews mechanism, worked example, and counter-positions for technical defensibility. Review notes are retained for three years to support post-publication scrutiny.
3. Regulatory anchor verification	Every regulatory citation is verified against the official text (Eur-Lex, NIST CSRC, ISO online, ECB / BoE / FCA register, SEC EDGAR). Article numbers and clause references are checked at the date of build.
4. Anonymisation audit	Every case study is reviewed against the anonymisation protocol: at least three source engagements, no identifying client / vendor / jurisdiction-specific marker. Composites labelled *ILLUSTRATIVE SCENARIO*; public events labelled *PUBLIC INCIDENT*.
5. Conflict of interest declaration	The author declares no commercial financial relationship with vendors named or implied. Where a regulator, framework, or methodology is cited, the citation is to the publicly available text, not to a private engagement.
6. Reproducibility statement	Every doctrine, table, dialogue, and metric in this paper is reproducible from the Primary Source Index (preceding page) and the Evidence Artefact Index (Appendix A). Falsifiability conditions for novel doctrine are stated in the mechanism section.

Editorial standard: If it cannot be evidenced, it cannot be defended. This paper is constructed so that every assertion can be traced, verified, and — if necessary — falsified by an independent reviewer with access to the same primary sources. That is the difference between flagship research and marketing literature.

TIER 1A · G L O S S A R Y

Glossary of Institutional Terms

Definitions below are paper-specific. Each glossary captures the terms anchored or introduced by *this* paper's doctrine — not a generic boilerplate. Where a term is the author's framework, it is marked with ™. Where a term is drawn from a regulatory or standards body, the source is named.

Term	Definition
Decision-Grade Reporting Standard™	Author framework: each cyber report is a decision frame, not a status update.
Decision Frame	A board-paper structure: one question, one recommendation, three options, cost / risk / time / evidence per option, signature block.
Board-Pack Pathology	Catalogue of failure modes in board cyber reporting; basis for the Upadrasta Board-Pack Pathology Study.
Senior Independent Director	Non-executive director designated as channel for shareholder concerns and challenge to the chair.
Disclosure Committee	Board sub-committee owning material-disclosure decisions, including SEC Item 1.05.
Risk Appetite Statement	Board-approved articulation of the level of risk the institution is willing to accept; cyber annex required.
Red-Amber-Green	Status-classification convention; necessary for status reports but insufficient for decision-grade reports.

TIER 1A · S C O P E

Limitations, Scope, and Defensibility Caveats

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

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

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

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

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

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

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

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

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

THE CLOSING DOCTRINE

The doctrine in one line.

The board's scarcest cyber resource is its own attention; the CISO's scarcest professional asset is the credibility to direct that attention. Both are spent on the wrong things by a 47-metric dashboard. Both are concentrated on the right things by four tier-1 metrics with bound decisions. The discipline is curation; the deliverable is governance; the residual is signed.

> ***"A dashboard reports what happened. A decision pack tells the board what to decide. The regulator credits only the second."***

Issued by: Kieran Upadrasta — CISSP · CISM · CRISC · CCSP · MBA · BEng
Affiliations: Schiphol University · Imperials · UCL · ISACA London (Platinum) · (ISC)² London (Gold) · PRMIA · ISF.
Contact: info@kieranupadrasta.com · www.kie.ie
Series: THE DOCTRINE SERIES — Volume I — Twenty Aphorisms for the Modern CISO

CLOSING APHORISM

"A dashboard reports what happened. A decision pack tells the board what to decide. The regulator credits only the second."

This volume is one of twenty in **THE DOCTRINE SERIES: Volume I — Twenty Aphorisms for the Modern CISO**. Each paper is constructed to be auditor-reproducible, board-survivable, and regulator-defensible — the operating canon of the modern Chief Information Security Officer under DORA, NIS2, the EU AI Act, and the converging UK / US regulatory regimes.

If it cannot be evidenced, it cannot be defended.

Kieran Upadrasta

CISSP · CISM · CRISC · CCSP · MBA · BEng

Cybersecurity Authority · Board Advisor · Interim CISO

www.kie.ie · info@kieranupadrasta.com · linkedin.com/in/kieranupadrasta

v4.1 ENGINEERING-INTEGRATED · CLOSING DOCTRINE

"In v4.0 we proved the engineering plane existed. In v4.1 we put it where it belongs — at the front of the doctrine, not the back. The Front Plate names the board question, the operating artefact, and the engineering. The artefact is screenshot-ready. The engineering is named and tooled. The v3.0 doctrine body is preserved — but now it is held up by the technical substrate that the supervisor, hiring manager, and procurement officer all need to see first."

Governance signs the doctrine. Engineering signs the deliverable.

v4.0 Engineering Plane closing aphorism — Doctrine Series Volume I.

If it cannot be evidenced, it cannot be defended.

Series umbrella aphorism — Doctrine Series Volume I.