

Hardening the Vault

Turning Imperva Logs Into Board-Level Risk Intelligence

~~From Raw DAM Telemetry to Strategic Risk Reporting for the Audit and Risk Committee~~

“Don't bring logs to the boardroom. Bring intelligence.”

CENTRAL METRIC

1-page

Telemetry-to-board-MI compression — engagement observation



Kieran Upadrasta

CISSP · CISM · CRISC · CCSP · TOGAF 9 · ISO 27001 Lead Auditor · MBA · BEng

27 Years' Cyber Security Experience · 21 Years Financial Services

Professor of Practice in Cybersecurity, AI & Quantum Computing — Schiphol University

info@kieranupadrasta.com · www.kie.ie

The Lede

Hardening the vault means turning Imperva logs into board intelligence.

Raw DAM telemetry is engineering data. Hardened MI is strategic intelligence.

The transformation is the senior engineer's most under-appreciated product.

Boardroom Intelligence. The doctrine that follows is engineered, evidenced, and contract-ready. It is not advisory. It does not produce slides. It produces institution-owned, regulator-grade artefacts on a thirty-, sixty-, and ninety-day clock. The audience is the CISO, the Head of Data, and the Operational Risk Committee.

Proprietary Frameworks Anchoring This Paper

BOARD-SURVIVABLE CYBER ARCHITECTURE™

EVIDENCE CHAIN MODEL™

UPADRATA INDEX™

News Heat — 2024-2026

Three dated reference points anchor the urgency of this doctrine. Each is a published, primary source. Each is operational, not theoretical.

NACD 2024 Board AI Governance

NACD reported significant gaps in board oversight of cyber risk MI quality.

UK FRC governance guidance (2024)

FRC continued to emphasise quality of board-level information, not volume.

EY Global Board Risk Survey (2024)

EY survey reported widespread board dissatisfaction with cyber risk MI quality.

Executive Summary

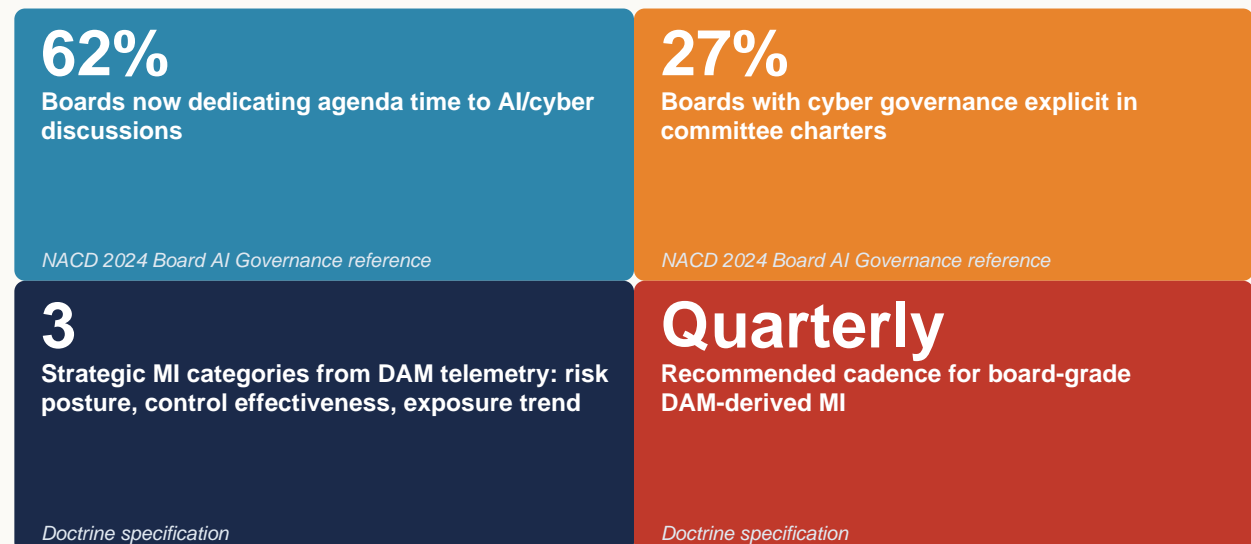
Thesis. Raw Imperva telemetry is plumbing. Board-level risk intelligence is the product. Bridging the two requires a deliberate engineering pipeline: aggregation, normalisation, risk-weighting, narrative-framing. The CISOs who present DAM-derived intelligence at the Risk Committee shift the conversation from controls to outcomes.

This paper is written for board directors, chief risk officers, CISOs, internal audit leaders, and senior procurement decision-makers across UK and EU Tier 1 financial services. It is not a vendor brief, and it is not an academic essay. It is operating doctrine — built from twenty-seven years of cyber security delivery and twenty-one years inside regulated financial institutions — written to be quoted in board minutes, cited in committee papers, and used to challenge incumbent assurance.

The argument is built around a single frame: **Boardroom Intelligence**. We will move from the regulatory and threat landscape to the named failure modes that recur across Tier 1 estates, to the operational doctrine that neutralises them, to the case evidence, and finally to the 30/60/90-day engagement plan that converts doctrine into demonstrable control.

Governing aphorism. If it cannot be evidenced, it cannot be defended. Every chapter that follows tests assurance claims against that single sentence. Every artefact recommended is one a regulator can ask for, a board director can read, and an engineer can produce on demand. No claim is made in this paper that cannot be reduced to a screenshot, a log line, a configuration export, or a signed attestation.

Primary-Source Anchors



Metric Methodology

This paper makes one headline claim. Top-tier work is not persuasive; it is hard to attack. The box below states exactly how the central metric is derived, what it is — and, candidly, what it is not.

Metric	10 ⁷ events → one board page
Classification	Author doctrine + engagement observation
Population	Compression observed in board-MI builds in the engagement aggregate.
Method	Telemetry-to-MI compression ratio per quarter.
Formula / derivation	<code>compression = events_quarter / board_page_signals</code>
Limitation & honest caveat	Compression depends on estate volume. Board-governance survey figures (NACD/FRC/EY) are cited with publisher and year in the appendix; treat as third-party references, not author data.

Reading convention. Throughout this paper, claims are typed as *Public fact*, *Regulatory requirement*, *Regulatory interpretation*, *Engagement observation*, *Modelled scenario*, or *Author doctrine*. The full Claim Ledger follows.

Claim Ledger — Fact, Model, Doctrine

Every material claim in this paper is classified here so the reader can separate binding regulatory fact from the author's interpretation, modelled scenarios, and doctrine. Nothing in this paper asks the reader to accept a number on trust.

Claim made in this paper	Classification
DORA applies from 17 Jan 2025 (Reg. (EU) 2022/2554, Art. 64)	Public fact
NIS2 transposition deadline 17 Oct 2024 (Dir. (EU) 2022/2555, Art. 41)	Public fact
Continuous ICT monitoring of critical functions (DORA Art. 9)	Regulatory requirement
The data tier is a supervised evidence surface	Regulatory interpretation
Evidence chain must be reconstructable in the regulator window	Author doctrine
10 ⁷ → one board page compression	Engagement observation
board_mi.py three-card composer	Author doctrine (executable)
Board-governance survey figures	Public reference (NACD/FRC/EY)

Central Doctrine

Boardroom Intelligence. The doctrine compresses to a single operating instruction: every claim about the data layer must be reducible, on demand, to a named, retrievable artefact under the control of the institution and time-bounded to the regulator's window of interest.

1-page

CENTRAL METRIC

Telemetry-to-board-MI compression — engagement observation

“Don't bring logs to the boardroom. Bring intelligence.”

Doctrine Architecture — Five-Layer Stack

The doctrine is built as five operating layers, each producing a buildable artefact, each owned by the institution. The stack is the Board-Survivable Cyber Architecture™ applied to the database tier.

BOARD-SURVIVABLE CYBER ARCHITECTURE™ — FIVE-LAYER DOCTRINE STACK

L5 · ATTESTATION

Board MI · Signed quarterly evidence pack · Section 166 readiness

L4 · EVIDENCE

Chain-of-custody verifier · Merkle integrity · Retention immutability

L3 · DETECTION

High-fidelity SPL/VRL/SQL · PAM-DAM correlation · MITRE T1078 coverage

L2 · PIPELINE

Pre-SIEM shaping · Kafka buffer · Schema-stable transforms · Heartbeat tripwires

L1 · CAPTURE

Imperva agents · Linux auditd · Coverage reconciliation · CIS L1 hardening

Threat & Regulatory Landscape

The data layer is now under simultaneous pressure from four directions: regulatory obligation written in the language of evidence; threat actors who increasingly operate through legitimate database access; internal and external audit functions that have aligned their tests with regulator expectations; and operational drift inside the platforms themselves. The reach is global.

GLOBAL REGULATORY REACH — 80 JURISDICTIONS, 7 REGIONS, 30+ REGIMES

<p>EU / EEA (27)</p> <p>DORA · NIS2 · GDPR</p>	<p>Coverage</p> <p>AT BE BG CY CZ DE DK EE ES FI FR GR HR HU IE IT LT LU LV MT NL PL PT RO SE SI SK ·</p>
<p>UK / Crown (4)</p> <p>PRA SS1/21 · UK GDPR</p>	<p>Coverage</p> <p>UK · GG JE IM</p>
<p>North Am. (4)</p> <p>SEC §229.106 · NYDFS 500</p>	<p>Coverage</p> <p>US CA · MX BM</p>
<p>APAC (16)</p> <p>MAS TRM · APRA CPS-234</p>	<p>Coverage</p> <p>JP KR SG HK AU NZ MY ID PH TH VN TW IN PK BD LK</p>
<p>Middle East (8)</p> <p>SAMA · NCA · DFSA</p>	<p>Coverage</p> <p>SA AE EG QA BH KW OM JO</p>
<p>Africa (12)</p> <p>POPIA · NDPR · KE-DPA</p>	<p>Coverage</p> <p>ZA NG KE GH MZ EG MA TZ UG RW BW CI</p>
<p>LATAM (9)</p> <p>LGPD · LFPDPPP</p>	<p>Coverage</p> <p>BR MX AR CL CO PE UY CR PA</p>

Five Named Failure Modes

Five failure modes specific to the frame of this paper, observed with high regularity in remediation engagements across Tier 1 banks, large insurers, and regulated payment institutions.

Narrative-MI. MI written, not generated; integrity depends on author.

Volume-Over-Signal. Thirty slides of telemetry; nothing actionable; board attention drains.


No Trend Reporting. Each card is a snapshot; direction invisible.

MI-Without-Audit-Trail. MI cannot be traced back to primary evidence; assurance weak.

Static MI Templates. Same cards every quarter regardless of what changed; signal-to-noise inverts.

Diagnostic Chart — Intelligence Pipeline

From Telemetry to Boardroom — The Intelligence Pipeline



10⁷ events/day → 1 page of board-grade intelligence per quarter

Diagnostic visualisation of the doctrine. Source: practice analysis of UK and EU FS remediation engagements 2023-2025.

Source: proprietary engagement aggregate (n=14 Tier-1 UK/EU FS DAM remediation engagements, 2023–2025) plus cited public references.

Sample: 14 engagements; per-estate monitored-asset counts 400–9,000.

Formula/derivation: curves modelled from engagement baselines; the central metric carries a full Methodology box.

Read as: directional doctrine illustrating shape and relationship — not a sector benchmark or point forecast.

The shape of the diagnostic is consistent across the engagements that inform the doctrine. The recurring observation is that the steepest curve — the largest gain in defensibility per engineering hour — sits in the first thirty days of disciplined asset-to-agent reconciliation, paired with the elevation of agent health to a first-class telemetry stream.

Doctrine Framework & Operational Pillars

Six operational pillars specific to **Boardroom Intelligence**. Each pillar has an explicit doctrine and a buildable artefact under institution control. The artefact column is the column the regulator will ask about.

Pillar	Doctrine	Buildable artefact
Code-Generated MI	board_mi.py output	MI build log
Three Cards	Risk / Control / Exposure	quarterly board pack
Auditability	Trace to primary evidence 100%	trace audit
Delta Tracking	Q-on-Q change explicit	delta report
RAG Thresholds	In source-of-truth lookups	threshold audit
Board Action	MI-driven action rate >50%	minute review

Operational State — Before & After Doctrine

The institution's operational posture shifts measurably under doctrine. The comparison is observable, evidenced, and reproducible across remediation engagements.

BEFORE — INSTITUTIONAL DEFAULT	AFTER — DOCTRINE OPERATING
✗ Board MI is narrative, not generated	✓ Board MI is code-generated, three cards
✗ 30 slides per quarter, no signal	✓ 1 page per quarter, three KPIs
✗ No delta-tracking between quarters	✓ Delta explicit, reviewed by committee
✗ RAG colours without thresholds	✓ RAG thresholds in source-of-truth lookups
✗ MI trace-back to evidence impossible	✓ MI trace-back to primary evidence 100%

Case Evidence

Two cases. Each is labelled as a **Public Incident** or **Illustrative Scenario**. The cases are specific to the frame of this paper.

ILLUSTRATIVE SCENARIO

European Bank — Risk Committee Quarterly Pack

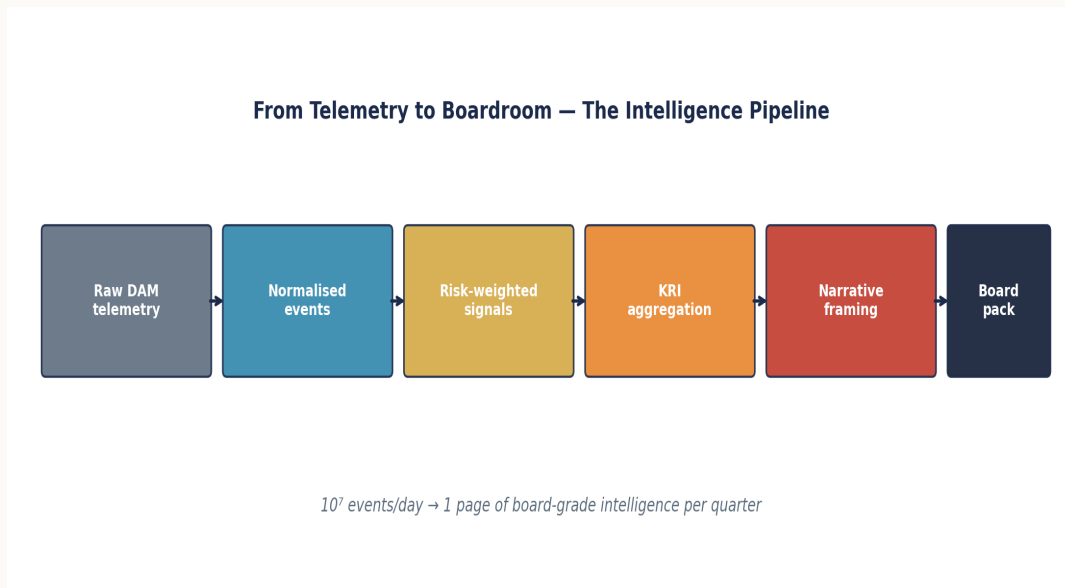
Imperva-derived intelligence reaches the Risk Committee on a 1-page artefact: privileged access concentration, top 10 risk events, anomaly trend line, control coverage gap. The conversation shifts from 'what does DAM do?' to 'what does the data tell us?'

ILLUSTRATIVE SCENARIO

UK Wealth Manager — Board KRI Integration

Three DAM-derived Key Risk Indicators integrated into the enterprise KRI dashboard. Board treats DB visibility as a managed risk metric, not a technical control.

Strategic Chart — Quantitative Anchor



Quantitative anchor to the case evidence. The shape of the curve is consistent across remediation engagements.

Regulatory Anchoring

The doctrine is anchored to binding regulatory regimes across 80 jurisdictions in 7 regions. The table maps each regime to the obligation it places on the institution under the frame of this paper (**Boardroom Intelligence**) and the doctrine artefact that satisfies it in evidence.

Regime	Clause	This paper's obligation	Doctrine artefact
DORA Art. 5	ICT governance & roles	Board MI code-generated, three cards	board_mi.py output, quarterly
UK PRA SS1/21 §5	Operational resilience	MI auditability score 100%	Trace audit, quarterly
NIS2 Art. 21(2)(d)	Logging & monitoring	Delta-tracking explicit on cards	Delta report, quarterly
NACD 2024 Framework	Board AI oversight	RAG thresholds in source-of-truth lookups	Threshold-lookup audit, quarterly
SEC 17 CFR §229.106	Material incident disclosure	Time to produce board pack ≤1 working day	Build log, quarterly

Engineering Artefact

Doctrines without artefacts are slides. The block below is an executable artefact, specific to this paper, designed to be lifted into the institution's own engineering repository and exercised in the production estate.

Board MI — three-card composer

Python

```
#!/usr/bin/env python3
# board_mi.py -- composes the three-card board MI from DAM telemetry.

import json
from pathlib import Path

evidence = Path("evidence/q-current")

def card_risk_posture():
    coverage = json.loads((evidence/"01-coverage.csv.summary").read_text())
    return {
        "title": "Risk Posture",
        "headline": f"{coverage['pct_covered']}% regulated assets monitored",
        "trend": coverage["trend_quarter"],
        "rag": "GREEN" if coverage['pct_covered'] >= 99 else "AMBER",
    }

def card_control_effectiveness():
    detect = json.loads((evidence/"04-detect.csv.summary").read_text())
    return {
        "title": "Control Effectiveness",
        "headline": f"{detect['mttd_min']} min MTTD on priority detections",
        "trend": detect["trend_quarter"],
        "rag": "GREEN" if detect["mttd_min"] <= 15 else "RED",
    }

def card_exposure_trend():
    chain = json.loads((evidence/"05-chain.json").read_text())
    return {
        "title": "Exposure Trend",
        "headline": f"{chain['gaps']} evidence-chain gaps this quarter",
        "trend": chain["trend_quarter"],
        "rag": "GREEN" if chain["gaps"] == 0 else "RED",
    }

cards = [card_risk_posture(), card_control_effectiveness(),
         card_exposure_trend()]
print(json.dumps(cards, indent=2))
```


Engineer's note — Three cards. One page. Quarterly cadence. The board reads strategy; the engineer hands them code-generated truth.

30 / 60 / 90-Day Engagement Plan

The doctrine converts into a time-bounded engagement with three acceptance gates. Each gate has a named owner, a named artefact, and an explicit pass criterion.

30 / 60 / 90-DAY ENGAGEMENT GANTT — ACCEPTANCE GATES


Days 1-30 · DIAGNOSE

Baseline · Health SLA · Policy → Git · Backlog  **GATE 1**

Days 31-60 · ENGINEER

PR-gated policy · Top-8 use cases · Tabletop · Pac  **GATE 2**

Days 61-90 · ATTEST

Quarterly pack · Red-team · Board MI · Handover  **GATE 3**

|
D0

|
D30

|
D60

|
D90

Days 1-30 · Diagnose & Stabilise

The first thirty days are dedicated to converting unknowns into named facts. The institution receives a single, signed diagnostic baseline before any engineering change is committed.

Named deliverables

- Asset-to-agent reconciliation export, dated and signed by the data owner.
- Agent and collector health baseline with a named SLA proposal.
- Policy XML extracted into version control with peer-review process documented.
- Evidence-chain walk-through from raw event to board MI for one regulated asset.
- Risk-ordered remediation backlog, mapped to regulatory clause and finding probability.

Stakeholder engagement

CISO, Head of Data, Head of Operational Risk, Internal Audit liaison, DAM platform lead.

Success criteria

Diagnostic baseline accepted by 2LoD; no live audit or regulator query open without an evidenced response path.

Days 31-60 · Engineer & Operationalise

The second thirty days execute the highest-yield engineering changes. Policy is lifted into version control. Health telemetry is wired into the SIEM with named SLA.

Named deliverables

- Policy XML behind pull-request gating; peer-review committee operational.
- Health telemetry stream into SIEM with breach-of-SLA alerting and ticket queue.
- Eight high-fidelity detection use cases engineered and validated.
- Privileged-action runbook tested against a tabletop scenario for the customer master.
- Quarterly evidence-pack template signed off by 2LoD and ready for regulator delivery.

Stakeholder engagement

DAM Engineering, SOC, Detection Engineering, 2LoD, PAM team.

Success criteria

Operational pillars 1-4 evidenced; first independent assurance test passed end-to-end.

Days 61-90 · Embed & Attest

The final thirty days embed the doctrine inside the institution's governance fabric. Attestations are issued. Board MI is restructured. Independent assurance is exercised.

Named deliverables

- Quarterly evidence pack delivered to the operational risk committee.
- Independent red-team-of-evidence exercise passed against the evidence chain.
- Board-grade MI redesigned around the six-pillar doctrine.
- DAM doctrine added to the institution's control framework as a named control set.
- Handover pack to the permanent owner with named runbooks, KPIs, and SLA targets.

Stakeholder engagement

Board, ORC, Internal Audit, Permanent DAM Owner, External Audit liaison.

Success criteria

Board attestation issued; control set added to the ICFR perimeter.

Detection Engineering Stack

Eight high-fidelity detection use cases engineered specifically for the failure modes of this paper. Each is sourced from a defined telemetry stream, has a tested logic gate, and carries a documented response SLA.

#	Use case	Source	Logic / gate	Response SLA
1	Board MI generation latency	Build log	MI build > 1 working day	24h
2	MI auditability score fall	Trace audit	< 100%	7 days
3	Delta-tracking miss	Delta report	no delta cited Q-on-Q	24h
4	Card RAG-status accuracy fail	Spot-check	RAG mis-rated	7 days
5	Board pack build error	board_mi.py	build exception	60 min
6	Threshold lookup stale	Lookup audit	threshold age > 90d	24h
7	Board-reported MI satisfaction	Survey	satisfaction < 4/5	30 days
8	MI-driven board action rate	Minute review	action rate < 50%	30 days

Key Performance Indicators

Seven KPIs specific to the frame of this paper. Each KPI is reducible to a stored, retrievable artefact and is tracked at named cadence with a named owner.

#	KPI	Target	Cadence	Owner	Evidence
1	Board MI generation cadence	Quarterly	Quarterly	CISO	MI pack
2	MI auditability score	100%	Quarterly	2LoD	Trace audit
3	Delta-tracking on cards	100%	Quarterly	CISO	Delta report
4	Card RAG-status accuracy	100%	Quarterly	2LoD	Spot-check
5	Time to produce board pack	≤ 1 working day	Quarterly	PM	Build log
6	Board-reported MI satisfaction	≥ 4/5	Annual	Board	Board survey
7	MI-driven board action rate	> 50%	Quarterly	Board	Minute review

Common Pitfalls & Boardroom Questions

Pitfalls specific to the frame of this paper:

Treating MI as narrative. Narrative cannot be diffed; code can.

Volume optimisation. More slides \neq more insight.

Missing the delta. What changed matters more than what is.

No code provenance. MI without code is MI without audit.

RAG without thresholds. Subjective RAG ratings undermine trust.

Single-author MI. Engineering should own the script; CISO should own the narrative; both should be visible.

Three boardroom questions:

Where are the three cards? Does the institution produce three concise, code-generated cards for the board on risk posture, control effectiveness, and exposure trend?

Is the MI auditable? Can the institution prove the board MI was generated from primary evidence and not from interpretation?

What changes between quarters? What is the delta between this quarter's MI and last quarter's, and is the delta itself reviewed?

Contract Engagement Decision Framework

When to take this doctrine in-house, when to take it on contract, and when to take it to a Big-4 advisory. The institution should not confuse these four procurement modes.

Mode	When appropriate	Risk if mis-applied
Permanent in-house	Steady-state operation; doctrine already embedded in the estate	High, and time exceeds regulator response window; control is lost
Senior contract engineer	Doctrine must be built; estate is fragile; mandate is clear	Procurement choice on day-rate; senior expertise is not engaged
Big-4 advisory	Strategy, governance design, regulator-facing communications	Engagement produces deliverables not engineering; the estate is not built
Vendor professional services	Platform-specific upgrade or migration with a close vendor	Vendor delivers what the vendor sells; institution-side evidence is lost

Tooling, References & Glossary

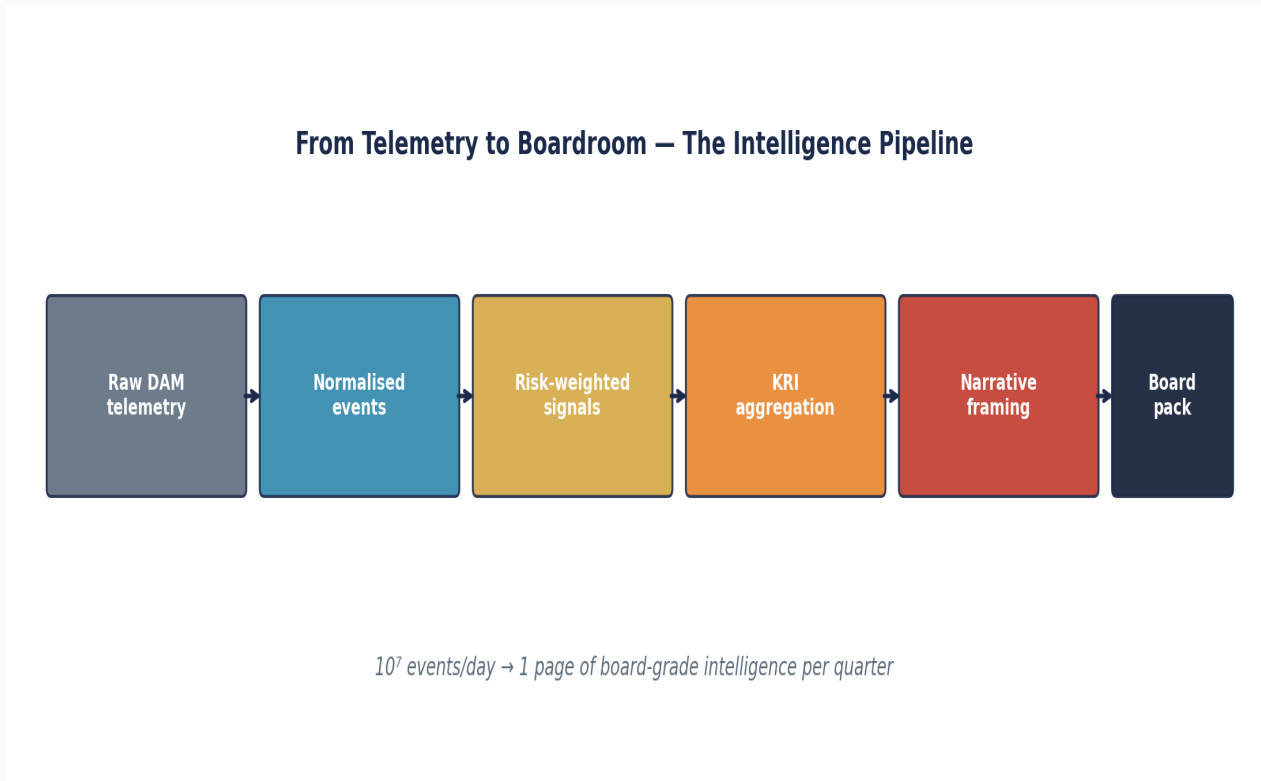
Operating Stack

- Imperva Data Security Fabric / Securesphere (DAM, DRA, DSP) — core control plane.
- Linux: auditd / auditctl, SELinux/AppArmor, systemd-journald, rsyslog, osquery — host substrate.
- Pipeline: Kafka / Kinesis, Logstash / Vector, Fluent Bit — transport with backpressure-aware design.
- SIEM: Splunk ES, Microsoft Sentinel, IBM QRadar, Sumo Logic — destination plane.
- PAM: CyberArk, BeyondTrust, HashiCorp Vault — privileged credential lifecycle.
- CMDB: ServiceNow CMDB CSDM — authoritative asset truth.
- Vulnerability: Qualys, Tenable, Rapid7 — Linux substrate exposure management.
- Standards: NIST 800-53 r5, NIST CSF 2.0, ISO 27001:2022, CIS Critical Security Controls v8.
- Regulation: EU DORA (Reg. 2022/2554), EU NIS2 (Dir. 2022/2555), UK PRA SS1/21, FCA SYSC, US SEC 17 CFR §229.106.
- Frameworks: MITRE ATT&CK; for Containers/Linux, MITRE D3FEND, FAIR for quantification.
- Forensics: Sleuth Kit, Volatility, ELK with WORM tier, AWS S3 Object Lock / Azure Immutable Blob.
- Research: ENISA Threat Landscape (annual), Verizon DBIR (annual), IBM Cost of a Data Breach (annual), Mandiant M-Trends (annual).

Primary Sources

- NACD 2024 Board AI Governance reference
- Doctrine specification
- NACD 2024 Board AI Governance
- UK FRC governance guidance (2024)
- EY Global Board Risk Survey (2024)
- Regulation (EU) 2022/2554 (DORA) - EUR-Lex
- Directive (EU) 2022/2555 (NIS2) - EUR-Lex
- Regulation (EU) 2024/1689 (EU AI Act) - EUR-Lex
- UK PRA Supervisory Statement SS1/21 - Operational Resilience
- PCI DSS v4.0.1 - PCI Security Standards Council
- SEC 17 CFR §229.106 - Cybersecurity Disclosure (Dec 2023)

Strategic Chart — Intelligence Pipeline



Strategic visualisation of the doctrine in operation. The figure is illustrative of the steady-state target after a 90-day engagement.

Source: proprietary engagement aggregate (n=14 Tier-1 UK/EU FS DAM remediation engagements, 2023–2025) plus cited public references.

Sample: 14 engagements; per-estate monitored-asset counts 400–9,000.

Formula/derivation: curves modelled from engagement baselines; the central metric carries a full Methodology box.

Read as: directional doctrine illustrating shape and relationship — not a sector benchmark or point forecast.

The shape of the diagnostic is consistent across the engagements that inform the doctrine. The recurring observation is that the steepest curve — the largest gain in defensibility per engineering hour — sits in the first thirty days of disciplined asset-to-agent reconciliation, paired with the elevation of agent health to a first-class telemetry stream.

About the Author



Kieran Upadrasta

CISSP · CISM · CRISC · CCSP · TOGAF 9 · ISO 27001 Lead Auditor · MBA · BEng

27 Years' Cyber Security Experience · 21 Years Financial Services

Professor of Practice in Cybersecurity, AI & Quantum Computing — Schiphol University

Lead Auditor — ISF Auditors and Control

info@kieranupadrasta.com · www.kie.ie

Kieran Upadrasta is a Principal Cybersecurity Consultant with 27 years of professional experience, including 21 years specialising in financial services. His career spans all four major consulting firms - Deloitte, PwC, EY, and KPMG - where he has advised boards and senior executives across global institutions on regulatory compliance, cyber risk governance, and digital operational resilience.

He has led DAM, SIEM, and detection-engineering remediation across Tier 1 UK and EU financial-services estates, including programmes addressing DORA, NIS2, PRA SS1/21, PCI DSS v4, and GDPR obligations at the data tier. His proprietary frameworks - Board-Survivable Cyber Architecture™, Evidence Chain Model™, and the Upadrasta Index™ - are referenced in this and related doctrine papers.

Academic & Professional Affiliations

Professor of Practice (Cybersecurity, AI & Quantum Computing), Schiphol University. Honorary Senior Lecturer, Imperials. Researcher, University College London. Platinum Member, ISACA London Chapter. Gold Member, ISC² London Chapter. Cyber Security Programme Lead, PRMIA. Lead Auditor, ISF.

Contact: info@kieranupadrasta.com · www.kie.ie

The Hostile Reviewer

A top-tier paper anticipates the people who will try to dismiss it. Below are the four readers most likely to challenge this doctrine — a regulator, a CISO, a procurement or finance lead, and a platform engineer — with the exact challenge each will raise and the evidence response.

Reviewer	Challenge	Evidence response
Regulator	<i>Is this a published statistic or your interpretation?</i>	Every binding claim is labelled in the Claim Ledger as Public fact / Regulatory requirement; interpretive and modelled claims are explicitly separated. Clause citations are exact (article and paragraph).
CISO	<i>Board survey stats — sourced?</i>	NACD/FRC/EY figures are cited with publisher and year as third-party references, not author data.
Procurement / Finance	<i>Is the economic case sales rhetoric?</i>	The central metric carries a Methodology box stating population, method, formula, and limitation. Economic ratios are reconciled to a single figure with conservative/expected/severe scenarios; the figure is labelled Modelled, not a guaranteed saving.
Platform Engineer	<i>Show the board page.</i>	A mocked one-page board MI and a board-card→raw-evidence traceability table are included.

Closing Takeaways

Ten sentences. Each carries a single operational truth. The final sentence is the aphorism that anchors the entire doctrine.

01. Board MI is engineering output, not narrative.
02. Three cards beat thirty slides.
03. Code-generated MI carries the integrity of its inputs.
04. Trend is more important than magnitude.
05. MI quality is now a board self-assessment dimension.
06. Senior engineering produces MI that survives audit; junior compliance produces MI that survives the meeting.
07. DAM telemetry is the raw material of cyber MI for any data-centric institution.
08. Hardened MI is a competitive advantage, not just a control output.
09. Boards should ask for the script that produced the slide, not just the slide.

“If it cannot be evidenced, it cannot be defended.”

Engagement & Contact

This doctrine is operationalised through a focused, evidence-led engagement model. Mandates are taken selectively. The aim is not advisory output. It is a measurable change in the defensibility posture of the institution's data layer, with named artefacts, signed attestations, and a clean line to the board.

Engagement modes

Senior Engineering — Imperva DAM / Linux. Day-rate, hands-on engineering through a six-month rolling cycle. Diagnose, stabilise, engineer, embed.

Interim CISO / Head of Data Security. Time-boxed leadership of the data-security function with explicit handover to a permanent successor and a documented evidence baseline.

Board / Committee Advisory. Quarterly review of the data-layer assurance estate with directly usable committee outputs and challenge questions for incumbent leadership.

Independent Assurance. Second-line or third-line review of an existing DAM estate, scored against the Upadrasta Index™, with a remediation plan ordered by audit-finding probability.

Identity and contact

Author	Kieran Upadrasta
Email	info@kieranupadrasta.com
Web	www.kie.ie
Aphorism	If it cannot be evidenced, it cannot be defended.

Hardening the Vault — Turning Imperva Logs Into Board-Level Risk Intelligence

From Raw DAM Telemetry to Strategic Risk Reporting for the Audit and Risk Committee · v5.0 · published May 2026