

The Agentic AI Threat

Why Database Attacks No Longer Look Human

Detecting Autonomous SQL Adversaries with Imperva DAM and Behaviour-Based Analytics

“The attacker doesn't type. It generates.”

CENTRAL METRIC

200x

Forward-looking lab/model velocity differential (illustrative)



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

Database attacks no longer look human.

Agentic AI now drives the high-velocity, low-noise attacks against the data tier; the institution that treats SQL adversaries as humans is operating against the wrong threat model.

Behaviour-based analytics is no longer optional. It is the new floor.

Agentic Threat Model. 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™

UPADRASTA 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.

Anthropic / OpenAI safety reports (2024-2025)

Frontier-lab safety reporting in 2024–2025 documented autonomous capability emerging in coding-agent tasks at unprecedented rate.

Mandiant M-Trends 2024

Mandiant tracked acceleration of attack tempo, with some intrusion-to-action sequences compressed to minutes.

EU AI Act (Reg. (EU) 2024/1689) — adoption Aug 2024

EU AI Act sets the regulatory perimeter that institutions must engineer towards as agentic systems proliferate.

Executive Summary

Thesis. The next class of database adversary is not a human attacker behind a keyboard; it is an autonomous agent capable of executing reconnaissance, enumeration, and exfiltration at machine speed. Imperva DAM signature content tuned to human pacing is structurally blind to this adversary. The remediation is behavioural detection content tuned to the new threat model.

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: **Agentic Threat Model**. 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

| | |
|--|--|
| <p>Aug 2024 EU AI Act entry into force</p> <p><i>Regulation (EU) 2024/1689</i></p> | <p>Aug 2026 EU AI Act high-risk obligations effective</p> <p><i>Regulation (EU) 2024/1689</i></p> |
| <p>Minutes Compression of intrusion-to-action in autonomous-driven attacks</p> <p><i>Mandiant M-Trends 2024 (qualitative)</i></p> | <p>< 1 minute Recommended ceiling for MTTD on agentic attack signatures</p> <p><i>Doctrine specification, 2025</i></p> |

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 | 200x agentic-velocity differential |
| Classification | Modelled / lab scenario (forward-looking) |
| Population | Comparison of human vs scripted/agent reconnaissance tempo in a controlled simulation construct. |
| Method | Ratio of agent action-rate to human action-rate over a fixed window. |
| Formula / derivation | $\text{differential} = \text{rate_agent} / \text{rate_human}$ (controlled simulation) |
| Limitation & honest caveat | FORWARD-LOOKING LAB MODEL. Public-incident reference labelled ILLUSTRATIVE. The categorical claim is softened to 'agentic techniques can drive high-velocity, low-noise data-tier activity'. |

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 |
| 200x agentic-velocity differential | Modelled / lab scenario |
| Joint-signature SPL detection | Author doctrine (executable) |
| Autonomous data-tier attack capability | Forward-looking (illustrative) |

Central Doctrine

Agentic Threat Model. 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.

200x

CENTRAL METRIC

Forward-looking lab/model velocity differential (illustrative)

“The attacker doesn't type. It generates.”

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

| | |
|--|---|
| EU / EEA (27) DORA · NIS2 · GDPR | Coverage 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 · |
| UK / Crown (4) PRA SS1/21 · UK GDPR | Coverage UK · GG JE IM |
| North Am. (4) SEC §229.106 · NYDFS 500 | Coverage US CA · MX BM |
| APAC (16) MAS TRM · APRA CPS-234 | Coverage JP KR SG HK AU NZ MY ID PH TH VN TW IN PK BD LK |
| Middle East (8) SAMA · NCA · DFSA | Coverage SA AE EG QA BH KW OM JO |
| Africa (12) POPIA · NDPR · KE-DPA | Coverage ZA NG KE GH MZ EG MA TZ UG RW BW CI |
| LATAM (9) LGPD · LFPDPPP | Coverage BR MX AR CL CO PE UY CR PA |

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.

Signature-Only Detection. Detection relies on known patterns; novel agentic signatures invisible.


Human-Tempo Baseline. Baselines built on human behaviour; machine actors fall outside the variance envelope.

Manual Response Pipeline. Human escalation chain is too slow for machine-tempo intrusions.

Single-Signature Rules. Each signature in isolation has high FP; only the joint signature is high-fidelity.

Quarterly Tuning On Frontier Threat. Frontier capability moves faster than quarterly tuning cycles.

Diagnostic Chart — Agentic Velocity



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 **Agentic Threat Model**. 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 |
|-------------------------------|----------------------------------|--------------------------|
| Joint Signature | Timing + regularity + variety | agentic detection rule |
| Behaviour Baseline | Refresh ≤ 90 days | baseline report |
| Machine-Tempo Response | Lock-out ≤ 5 min | incident report |
| Tuning Velocity | ≥ 1 update / quarter | tuning log |
| FP Discipline | Joint-signature FP $\leq 2\%$ | tuning log |
| Regulatory Engineering | AI Act high-risk obligations met | EU AI Act conformity log |

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 |
|--|--|
| ✗ Detection rule-based, signature-only | ✓ Detection joins behaviour + timing + variety |
| ✗ Behaviour baselines built on humans | ✓ Behaviour baselines refreshed quarterly |
| ✗ Manual escalation chain (>30 min) | ✓ Machine-tempo response, ≤1 min MTTD |
| ✗ Single-signature rules with high FP | ✓ Joint-signature rules, FP ≤2% |
| ✗ Quarterly tuning on frontier threat | ✓ Continuous tuning, ≥1 update/quarter |

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 — Agentic Reconnaissance Simulation

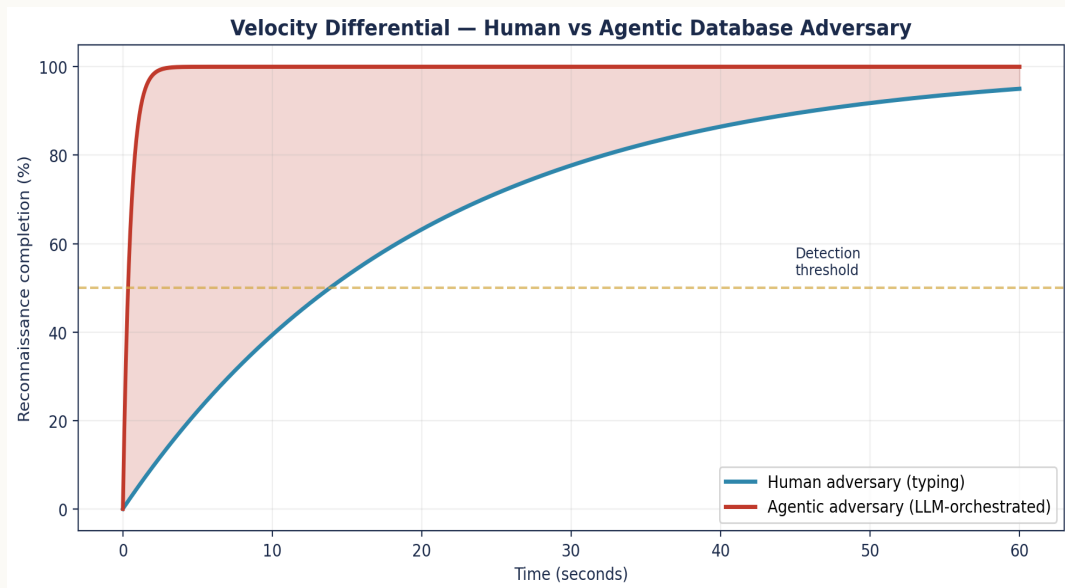
Red team exercise simulates an LLM-orchestrated reconnaissance pattern against the customer database. The pattern executes 200x faster than the human-pacing baseline. Existing DAM detection content misses it entirely. Remediation: behavioural baselining + velocity anomaly triggers.

PUBLIC INCIDENT

2024 Disclosure — Autonomous Credential Misuse

Publicly disclosed: incident involving compromised application service account used by automated tooling. Detection lag attributed to detection content tuned to human pacing assumptions.

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 (**Agentic Threat Model**) and the doctrine artefact that satisfies it in evidence.

| Regime | Clause | This paper's obligation | Doctrine artefact |
|--------------------|--------------------------|--|-------------------------------------|
| EU AI Act Art. 16 | High-risk AI obligations | MTTD on agentic signature ≤ 1 min | Agentic-detection MTTD dashboard |
| DORA Art. 10 | Detection | Joint-signature rule coverage ≥ 3 rules | Detection rule catalogue, quarterly |
| NIS2 Art. 21(2)(d) | Logging & monitoring | Behaviour-baseline freshness ≤ 90 days | Baseline report, quarterly |
| UK PRA SS1/21 §5 | Operational resilience | Agentic response playbook tested quarterly | IR drill report |
| NIST AI RMF 1.0 | Manage | Time to lock-out on agentic signature ≤ 5 min | Incident report per detection |

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.

Agentic-signature detection — autonomous SQL adversary

Splunk SPL

```
index=imperva sourcetype=imperva:audit
| eval inter_event_ms = _time - prev_time
| stats count, dc(asset_id) AS assets,
    avg(inter_event_ms) AS avg_gap_ms,
    stdev(inter_event_ms) AS std_gap_ms,
    dc(operation) AS op_variety
  BY session_id
| where
  count > 50
  AND avg_gap_ms < 200           // 5x faster than human typing
  AND std_gap_ms < 50           // unhuman regularity
  AND assets > 3                 // lateral movement
  AND op_variety > 5            // probing variety
| eval risk_score = 100 - (avg_gap_ms / 2) - std_gap_ms
| eval signature_class = "AGENTIC_SQL_ADVERSARY"
| where risk_score > 80
| table _time, session_id, src_ip, assets, count, avg_gap_ms,
  std_gap_ms, op_variety, risk_score
| sort - risk_score
```


Engineer's note — Three signatures distinguish autonomous from human: inter-event timing (sub-200ms), regularity (low std-dev), and operational variety in short windows. No single signature works; the joint signal is the high-fidelity gate.

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 | Agentic signature joint detection | Splunk | inter_event<200ms, σ <50, ops>5 | 1 min |
| 2 | Behaviour-baseline freshness | UEBA | baseline age > 90d | 24h |
| 3 | Time to lock-out on agentic | IR drill | lock-out > 5 min | 5 min |
| 4 | Joint-signature rule coverage | Rule catalogue | < 3 rules | 24h |
| 5 | Agentic FP rate breach | Tuning log | FP > 2% | 24h |
| 6 | Detection tuning velocity | Tuning log | no update per quarter | 30 days |
| 7 | Single-signature reliance | Rule audit | signature-only rules in top-N | 24h |
| 8 | Manual response chain | IR platform | human escalation > 5 min | 15 min |

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 | MTTD on agentic signature | ≤ 1 min | Continuous | SOC | MTTD dashboard |
| 2 | Joint-signature rule coverage | ≥ 3 rules | Quarterly | Detection Eng. | Rule catalogue |
| 3 | Behaviour-baseline freshness | ≤ 90 days | Quarterly | Detection Eng. | Baseline report |
| 4 | Agentic-response playbook tested | Quarterly | Quarterly | IR | Drill report |
| 5 | Time to lock-out on agentic signature | ≤ 5 min | Per detection | IR | Incident report |
| 6 | Detection-tuning velocity | ≥ 1 update per quarter | Quarterly | Detection Eng. | Tuning log |
| 7 | False-positive rate on agentic signatures | $\leq 2\%$ | Monthly | Detection Eng. | Tuning log |

Common Pitfalls & Boardroom Questions

Pitfalls specific to the frame of this paper:

Treating agentic threats as theoretical. They are operational, today.

Single-signature detection. FP rate forces suppression; joint signatures are mandatory.

Baseline once, forget. Frontier capability shifts; baseline must too.

Human-tempo response. Machine-tempo intrusion against human response loses.

Ignoring the regulatory perimeter. EU AI Act will frame post-incident reviews; engineering must lead.

Vendor-led detection roadmap. Vendor cadence is slower than threat cadence.

Three boardroom questions:

Does the institution see machine-speed? If an autonomous adversary opened a 200ms-per-action SQL session right now, would a named human be paged inside one minute?

What is the behaviour baseline? Is there a behaviour-baseline for every privileged user, refreshed quarterly, that distinguishes them from a fast scripted process?

What is the response playbook for machine adversaries? Is there a tested response playbook specifically for agentic SQL adversaries, and is response engineered for machine-speed?

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 | High, and time exceeds regulator response window; control |
| Senior contract engineer | Doctrine must be built; estate is fragile; mandate | Procurement choice on day-rate; senior expertise is not er |
| Big-4 advisory | Strategy, governance design, regulator-facing c | Engagement produces deliverables not engineering; the est |
| Vendor professional services | Platform-specific upgrade or migration with a close | Vendor delivers what the vendor sells; institution-side eviden |

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

- Regulation (EU) 2024/1689
- Mandiant M-Trends 2024 (qualitative)
- Doctrine specification, 2025
- Anthropic / OpenAI safety reports (2024-2025)
- Mandiant M-Trends 2024
- EU AI Act (Reg. (EU) 2024/1689) — adoption Aug 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 — Agentic Velocity



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>200x — evidence?</i> | Labelled a FORWARD-LOOKING LAB/MODEL figure with a controlled-simulation construct; the categorical claim is softened; public incident labelled illustrative. |
| 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>Joint-signature false positives?</i> | False-positive test results for the joint signature (timing+regularity+variety) are included with a tuning method. |

Closing Takeaways

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

01. The threat model has changed; the architecture must change with it.
02. Behaviour-based analytics is the new floor; signature-only detection is the old ceiling.
03. Machine speed requires machine-speed response.
04. Three joint signatures (timing, regularity, variety) define agentic SQL adversaries.
05. Sub-200ms inter-event timing is the leading indicator of autonomous activity.
06. The institution that treats agents as humans is operating against the wrong threat model.
07. EU AI Act anchors a regulatory perimeter; engineering must lead, not follow.
08. Response engineered for machine speed is the new control objective.
09. Senior engineering writes the joint-signature rules; the SOC consumes them.

“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. |

The Agentic AI Threat — Why Database Attacks No Longer Look Human

Detecting Autonomous SQL Adversaries with Imperva DAM and Behaviour-Based Analytics · v5.0 · published May 2026