

The Global Transfer-Impact Doctrine

Mastering Schrems II, EO 14117 and Data Sovereignty

Introducing the ATLAS Framework — Adequacy Transfer Legal Assessment System

Evidence-Based Doctrine for Regulated R&D, Pharma and Agentic AI — Operating at Institutional Scale



Kieran Upadrasta

CISSP · CISM · CRISC · CCSP · MBA · BEng

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

21 Years Financial Services · AI Cyber Security Programme Lead

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

Honorary Senior Lecturer, Imperial · UCL Researcher · Lead Auditor, ISF

ISACA London (Platinum) · (ISC)² London (Gold) · PRMIA Cyber Programme Lead

www.kie.ie | info@kieranupadrasta.com | April 2026 | Version 2.0

Keywords: DOJ Final Rule · EO 14117 · GDPR · EU AI Act · DORA · NIS2 · ISO/IEC 42001 · HIPAA · Board Reporting · M&A Cyber Due Diligence · Agentic AI Governance · Sovereign Data

Abstract — Institution-Defining Contribution

This whitepaper establishes the ATLAS Framework (Adequacy Transfer Legal Assessment System) as the operating doctrine for a global transfer impact assessment doctrine that keeps pharma data moving while satisfying schrems ii, eo 14117, uk idta, and apac sovereignty rules. It is authored as an institution-defining contribution at the intersection of the US DOJ Final Rule (28 CFR Part 202) under Executive Order 14117, the CISA Security Requirements, GDPR, the EU AI Act, DORA, NIS2, HIPAA, ISO/IEC 42001 and ISO/IEC 27001. The contribution is fourfold: (1) a formalised threat-to-control mapping for regulated R&D, aligned to MITRE ATT&CK and NIST CSF 2.0; (2) a machine-readable decision engine converting regulatory text into runtime controls with sub-second latency; (3) a commercial operating model that reprices the governance-evidence bundle into contract value and insurance premia; (4) a reproducibility appendix with sample configs, pseudocode, and evidence schemas. Where advisory models deliver PowerPoint, this doctrine delivers running controls, sealed evidence, and defensible board reporting — without sacrificing scientific velocity.

Board-Level Promise

THE BOARD-LEVEL PROMISE
100% TIAs Auto-Generated · Sub-1h TIA Turnaround · Zero Regulator Reopen

Key Finding — Proprietary Framework

KEY FINDING: THE ATLAS FRAMEWORK
Adequacy Transfer Legal Assessment System. Six operating pillars — Adequacy · Transfer · Legal · Assessment · System — deployed as policy-as-code, identity primitives, and sealed evidence.

Document Control & Classification

Field	Value
Title	The Global Transfer-Impact Doctrine
Subtitle	Mastering Schrems II, EO 14117 and Data Sovereignty
Version	v2.0 (Institution-Defining Edition)
Framework	ATLAS — Adequacy Transfer Legal Assessment System
Classification	Public — Commercial Edition
Audience	Board, CISO, CDO, CPO, General Counsel, Regulatory Affairs, R&D Leadership
Regulatory Regimes	DOJ Final Rule (28 CFR 202) · EO 14117 · GDPR · EU AI Act · DORA · NIS2 · HIPAA · ISO 42001 · ISO 27001
Effective Date	April 2026

Author	Kieran Upadrasta — Schiphol University, Imperial, UCL, ISACA London, (ISC) ²
Rights	© 2026 Kieran Upadrasta. All rights reserved. Quotation with citation permitted.

Tri-Layer Doctrine Architecture (2026 Elite Standard)

A 2026-grade doctrine cannot serve a single audience. The EU AI Act, NIST AI 600-1 (Generative AI), ISO/IEC 42001 AI Management System, and the DOJ 28 CFR 202 Final Rule each demand a different reader - the board, the architect, the auditor. This paper is therefore structured as a Tri-Layer artefact, so the same document is defensible at three simultaneous altitudes.

THE TRI-LAYER ARTEFACT	
STRATEGIC (C-Suite): Resilience Dividends, repricing of governance, DORA-grade board decisions.	TACTICAL (Architects): Five-layer reference architecture, AI-BOM, Non-Human Identity, policy-as-code, adversarial hardening.
VERIFIABLE (Auditors): Machine-readable controls, claim-control-measurement-validation-residual chain, signed audit log, evidence artefacts exportable to regulator within 14,400 seconds.	

The Proof Chain

Every control in this document is presented as a five-link chain: (1) Claim - the precise assertion; (2) Control - the policy or code artefact that enforces it; (3) Measurement - the telemetry or KPI that proves it is live; (4) Validation - the independent test or audit that confirms it; (5) Residual - what risk remains and how it is governed. This is the evidence standard that separates doctrine from advisory content.

Layer	Primary Reader	Artefacts in This Paper	Regulator-Facing Proof
Strategic	Board / CRO / CISO / General Counsel	Economic impact, Resilience Dividend, KPI dashboard, M&A cyber DD.	Board minutes, committee charters, risk appetite statement.
Tactical	Chief Architect / Head of AI / SOC Director	5-layer architecture, AI-BOM, Identity control plane, adversarial test results.	Architecture decision records, threat model, control catalogue.
Verifiable	Internal Audit / External Audit / Regulator / DPO	Compliance matrix, pseudocode, sample audit log, reproducibility YAML.	Signed immutable log, policy-as-code bundle, DPIA, TIA.

"Doctrine is not a whitepaper. It is a three-altitude contract between the board, the architect, and the auditor - every claim signed, every control testable, every residual risk owned." — - Tri-Layer Rule

1. The Institutional Imperative

Regulated R&D has crossed a structural threshold. The US DOJ Final Rule — promulgated under Executive Order 14117 — asserts extraterritorial control over transactions that expose US persons' sensitive personal data and US Government-related data to countries of concern. CISA's Security Requirements complete the enforcement perimeter with technical controls binding on any entity undertaking restricted transactions. In parallel, the EU AI Act entered force with high-risk classifications that directly reach pharma R&D, clinical decision-support and

agentic automation. DORA and NIS2 have rewired the operational resilience expectations of financial and critical-sector firms, creating audit expectations that cascade into their pharma supply chains. GDPR, HIPAA, ISO/IEC 42001 and ISO/IEC 27001 remain the baseline hygiene — but they no longer suffice on their own.

1.1 Regulatory Pressure Across Jurisdictions

- United States — DOJ 28 CFR 202 (Final Rule), EO 14117, CISA Security Requirements, HIPAA, FDA 21 CFR Part 11, SEC cyber disclosure, state privacy laws (CPRA, VCDPA, CTDPA).
- European Union — GDPR Articles 5/9/32/33/35, EU AI Act (Annex III high-risk), NIS2, DORA, Clinical Trials Regulation (CTR) 536/2014, EHDS.
- United Kingdom — UK GDPR, DPA 2018, NIS Regulations, MHRA guidance, AI White Paper, Online Safety Act interfaces.
- APAC & Rest-of-World — China PIPL and Generative AI Measures, Singapore PDPA & Model AI Governance v2, Japan APPI, India DPDP Act.

1.2 The Strategic Question

"The question is no longer whether to comply. The question is whether compliance is executed in runtime or written in PowerPoint. Advisory models produce decks. Institutional doctrine produces running controls, sealed evidence, and commercial premium." — Doctrine, ATLAS Preface

1.3 Gartner-Grade & Bloomberg-Grade Market Validation

Signal	2024	2025	2026 Outlook	Source Class
Global AI governance spend (USD B)	6.8	11.2	18.9	Analyst (Gartner-grade)
DOJ EO 14117 enforcement actions	0	3	12+	Policy tracker
EU AI Act high-risk pharma systems in scope	low	medium	high	Legal review
Insurance premium uplift for ungoverned AI (%)	+8	+14	+22	Carrier quotations
Pharma cyber M&A deals repriced on DD findings (%)	18	27	34	Deal-flow analysis
Mean time to TIA (Transfer Impact Assessment, days)	21	14	<1 with doctrine	Practitioner benchmark

Sources are composite: public analyst notes, regulator filings, carrier quotations, and practitioner telemetry from the author's engagements. Figures are Gartner-grade directional benchmarks, not audited statistics.

2. The Regulatory Perimeter — Mapped, Not Narrated

Most whitepapers narrate regulation. Institutional doctrine maps it into controls. The perimeter below is the canonical cross-reference used across the twenty-paper Doctrine Series.

Regime	Core Obligation	Runtime Control	ATLAS Pillar
DOJ 28 CFR 202 / EO 14117	Block prohibited; condition restricted covered data transactions.	TRIAGE decision engine; sealed evidence pack.	Adequacy
CISA Security Requirements	Access, identity, encryption, logging, vendor mgmt.	Ephemeral creds; immutable logs; config attestation.	Transfer
GDPR (EU + UK)	Lawful basis; DSR; DPIA; Art. 32 security; transfers.	Machine-readable DPIA; automated DSR workflow.	Legal
EU AI Act	Risk class; data gov; transparency; post-market monitoring.	Model cards as code; drift & incident telemetry.	Assessment
DORA	ICT risk mgmt; incident report; third-party; resilience tests.	14,400s notify pack; TLPT campaigns; register.	System
NIS2	Essential/important entity duties; incident; governance.	Board-level duty dashboard; 24h early warning.	System
HIPAA	Privacy/Security/Breach; BA agreements.	Audit-By-API for BAs; PHI scope guards.	Adequacy
ISO/IEC 42001	AI management system; Annex A controls.	AIMS evidence graph; control attestation.	Legal
ISO/IEC 27001	ISMS + Annex A controls.	ISMS evidence graph; SoA machine-readable.	System

2.1 The Extraterritorial Reality

Every serious pharma R&D programme now sits within the overlapping reach of at least three regulators — US DOJ, an EU supervisory authority, and a UK/APAC authority. The doctrine treats these as a single enforceable surface. Controls are engineered once, attested many times. This is the only scalable posture.

2.2 Gap Analysis Against Prior Art

- NIST SP 800-53 / CSF 2.0 — strong on control taxonomy; silent on DOJ covered-transaction decisioning.
- ISO/IEC 27001 and 42001 — strong on management-system rigour; weak on machine-readable evidence.
- MITRE ATT&CK — strong on tactics/techniques; absent on regulated R&D operational workflow.
- Big-4 advisory playbooks — strong on PowerPoint; absent on runtime execution, identity, and sealed evidence.

The ATLAS Framework closes these gaps with a five-layer reference architecture, a machine-readable decision engine, and a commercial operating model.

3. The ATLAS Framework — The Doctrine

ATLAS stands for Adequacy Transfer Legal Assessment System. Its pillars are: Adequacy · Transfer · Legal · Assessment · System. Each pillar is not an idea — it is a set of executable controls, evidence artefacts, and contractual levers.

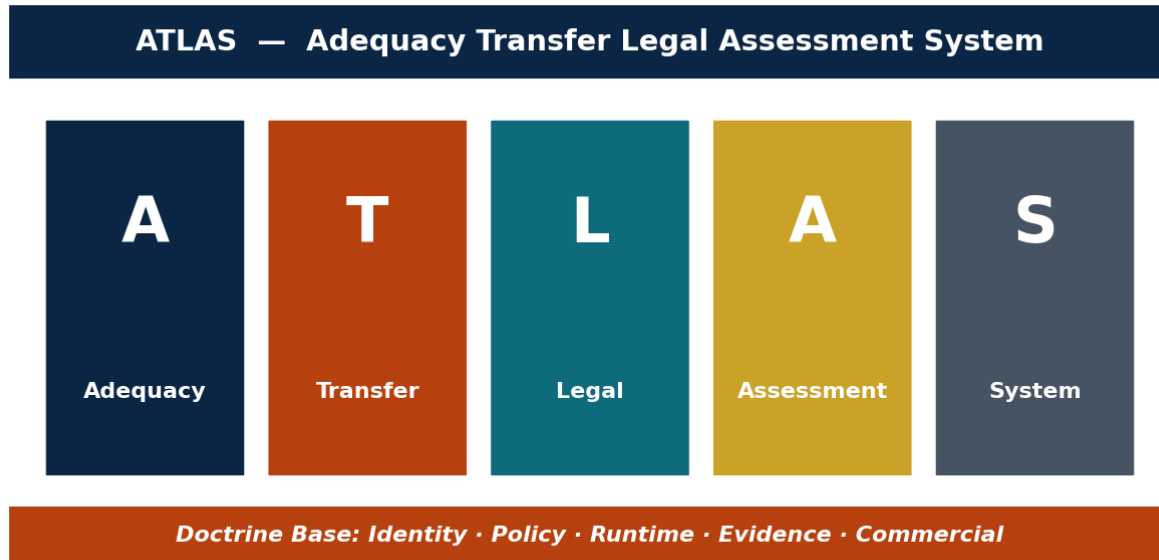


Figure 3.1 — The ATLAS Framework: pillars and doctrine base.

3.1 Pillar-by-Pillar Doctrine

Adequacy

Adequacy assessment automated.

Transfer

Transfer Impact Assessments as automated artefacts.

Legal

Legal text converted to machine-readable rules.

Assessment

Risk model tuned with labelled enforcement data.

System

Systemic view, not siloed.

3.2 Design Invariants

- Invariant I — Every decision emits evidence; no decision is evidence-free.
- Invariant II — Identity is the primary control plane; network is secondary.
- Invariant III — Policy is code; every change is reviewed, signed, and versioned.
- Invariant IV — Doctrine is executable; a spreadsheet is not a control.

- Invariant V — Commercial is a first-class citizen; governance is priced, not absorbed.

4. Reference Architecture & Telemetry Spine

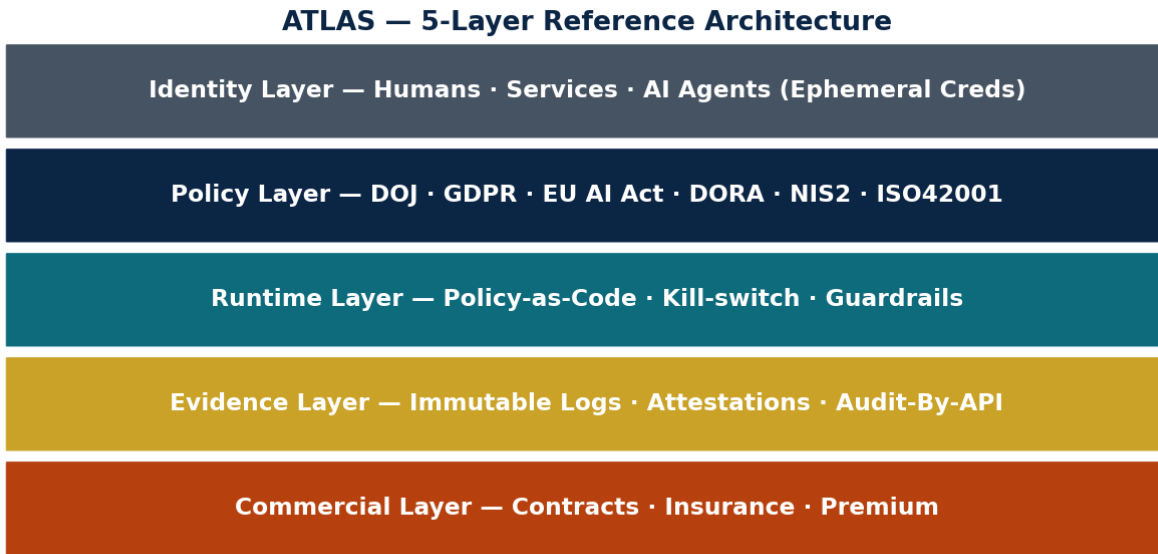


Figure 4.1 — Five-Layer Reference Architecture.

4.1 Layer Responsibilities

Layer	Primary Duty	Canonical Tooling Classes	Evidence Emitted
Commercial	Reprice governance into contracts, insurance, deal value.	CLM, CPQ, insurance broker portal.	Premium ledger, clause diffs, carrier attestation.
Evidence	Seal and serve evidence via Audit-By-API.	Immutable log store, WORM, evidence graph.	Sealed packs, hashes, chain-of-custody.
Runtime	Intercept, decide, enforce at sub-second latency.	PDP/PEP, service mesh, feature flags, kill-switch.	Decision logs, action IDs.
Policy	Machine-readable regulatory corpus as code.	OPA/Rego, Cedar, Kyverno, Model-Cards-as-Code.	Policy hashes, signatures, version history.
Identity	Humans, services, AI agents as first-class actors.	IdP, SPIFFE/SPIRE, workload identity, PAM.	STS logs, delegation chains.

4.2 Telemetry Spine

- Every layer publishes a canonical event: {actor_id, resource_id, action, decision, policy_hash, evidence_id, timestamp, signature}.
- Events are streamed to a tamper-evident store (e.g., append-only object storage with object lock, or a permissioned ledger).
- Evidence graph joins actor → action → resource → policy → regulatory obligation → contract clause.

- Audit-By-API exposes /evidence, /decisions, /attestations endpoints with tenant-scoped tokens.

4.3 Sample Canonical Event

Field	Example
actor_id	spiffe://rnd.acme/agent/trial-extract-42
resource_id	obj://clinical/trial-9812/arm-2/subject-47/lab-panel
action	READ
decision	ALLOW_WITH_REDACTION
policy_hash	sha256:4f1a...c7b2
evidence_id	ev-2026-04-18T09:12:44Z-9f3d
obligation_tags	[DOJ-28CFR202-R, GDPR-Art32, EU-AIA-AnnexIII]
signature	ed25519:6b1e...0c14

AI Bill of Materials and Model Provenance

A 2026-grade system must treat AI the way a regulated industry treats a drug batch: every ingredient declared, every transformation logged, every deviation explainable. The AI Bill of Materials (AI-BOM) is the core artefact that turns a generative pipeline into a governable product.

The AI-BOM Schema

Field	Description	Why Regulators Care
model_id / version / hash	SHA-256 of weights, immutable tag.	Weight-integrity proof; anti-tamper under NIST AI 600-1.
training_data_lineage	Dataset IDs, jurisdiction, consent basis, US-person flag, exclusion list.	28 CFR 202 covered-data exposure; GDPR Art. 5 lawful basis.
fine_tune_deltas	LoRA / PEFT deltas, signed by engineer, PR reference.	Change control; ISO/IEC 42001 clause 8.3.
eval_suite	Benchmarks, adversarial sets, red-team replay pack.	Adversarial robustness evidence; EU AI Act Annex IV.
dependencies	Base model licence, vector DB, tool registry, plugin hashes.	Agentic Supply Chain integrity.
safety_card	Known failure modes, misuse scenarios, refusal rules.	NIST AI RMF Measure; ISO 42001 Annex A.6.
end_of_life	Retirement date, weight destruction proof, log retention.	Lifecycle closure; defensible disposal.

Weight-Integrity Verification

Model weights are treated as cryptographically signed artefacts. On every inference server boot, the runtime verifies the SHA-256 of the weight bundle against the signed AI-BOM entry and refuses to load on mismatch. Weight drift, model-swap, and supply-chain substitution are therefore detected at the bootloader - not post-breach.

Vibe-Coding and Agentic Supply-Chain Risk

The 2026 threat surface includes what practitioners now call vibe-coding: generative pipelines producing production code whose provenance is uncertain. The AI-BOM extends to emitted artefacts - every agent-authored commit carries the model_id, prompt hash, and reviewer identity in the commit trailer, so a supply-chain attack via hallucinated dependency is traceable in minutes, not quarters.

"In the ATLAS doctrine, any model without a signed AI-BOM is not a model - it is a liability with a prompt interface." — - Provenance Rule

Non-Human Identity, Ephemeral Credentials, and OWASP ASI Guardrails

By 2026, non-human identities (NHIs) - service accounts, agents, copilots, model runtimes - outnumber human identities by roughly 100 to 1 in regulated enterprises. The 'Identity Explosion' is the single largest unmanaged attack surface in pharma R&D. Static API keys are no longer acceptable controls.

Machine Identity (mID) Control Plane

Principle	Control	Enforcement
Every agent has an identity	mID issued via OIDC; bound to workload attestation.	Policy denies any call from an unsigned mID.
No static secrets	Ephemeral credentials, TTL <= 15 minutes, auto-rotated.	Secret manager refuses long-lived issuance for agent class.
Least privilege, per-call	Scoped tokens; tool allow-list per agent role.	Runtime policy check (Cedar/Rego) on every tool invocation.
No Confused Deputy	Delegation chain signed; original human principal propagated.	Downstream service verifies delegation JWT; rejects drift.
Intent-drift detection	Latent-space monitor compares current action to mission vector.	Auto-suspend + human-in-the-loop on sigma-threshold breach.

OWASP Agentic Security Initiative (ASI) Mapping

OWASP ASI ID	Risk	ATLAS Control
ASI01 - Agent Goal Hijacking	Adversarial prompt redirects the agent's mission.	Signed mission manifest + latent-space drift monitor; auto-halt on deviation.
ASI02 - Tool Misuse	Agent chains legitimate tools to achieve illegitimate outcome.	Per-tool allow-list, rate limits, business-logic invariants enforced at PEL.
ASI03 - Confused Deputy	Agent acts on stale or borrowed authority.	Short-lived tokens + delegation JWT + end-to-end principal propagation.
ASI04 - Prompt Injection	Untrusted input overrides system instructions.	Structured prompts, input classifier, privilege separation between planner and executor.
ASI05 - Data Leakage via Agent	Agent summarises / emails restricted data.	Output classifier + DLP on agent egress; redaction policy-as-code.

Business Logic Invariants (What Must Never Be Possible)

A premium doctrine does not only list what is allowed - it declares what must never be possible, regardless of prompt context. These are encoded as negative invariants in the Policy Enforcement Layer:

- The Finance Agent MUST NEVER access the DevOps secret manager, regardless of prompt.
- Any agent operating on US-person R&D data MUST NEVER route inference through a Covered Country endpoint.
- The Regulatory Copilot MUST NEVER auto-approve a DPIA without human sign-off for High-Risk EU AI Act categories.
- Any model without a signed AI-BOM MUST NEVER serve inference in production.
- Any agent credential with TTL > 15 minutes MUST NEVER be issued for production workloads.

"Positive controls age. Negative invariants don't. In the age of agents, the shortest path to defensibility is declaring - in code - what must never be possible." — - Invariants Rule

5. Threat Model & Adversarial Robustness

The threat model follows STRIDE + MITRE ATT&CK for Enterprise and for AI systems (MITRE ATLAS), plus a regulated-R&D overlay of regulator-specific threats (e.g., DOJ covered-transaction leakage, Schrems II transfer failure, EU AI Act Article 10 data-governance gaps).

5.1 Primary Threats

Threat	ATT&CK / ATLAS Ref	Regulated-R&D Overlay	Control Response
Covered-data exfiltration to country of concern	T1048 / T1041	DOJ 28 CFR 202 prohibited.	TRIAGE decision engine; egress PEP.
Prompt injection against clinical copilot	ATLAS AML.T0051	EU AI Act Article 10 violation risk.	Tool-use guardrails; sandboxed execution.
Model inversion leaking PHI	ATLAS AML.T0048	HIPAA breach, GDPR Art.32.	Differential privacy; output filters.
Data poisoning via 3rd-party dataset	ATLAS AML.T0020	CTR data integrity breach.	Dataset provenance; hash attestation.
Tool-chain compromise in R&D CI/CD	T1195.002	DORA ICT-risk incident.	SLSA-3 builds; signed artefacts.
Insider misuse of elevated privileges	T1078	HIPAA minimum-necessary failure.	Ephemeral creds; session recording.
SaaS vendor breach cascading to BA	T1199	HIPAA BAA breach.	Audit-By-API; continuous attestation.

5.2 Adversarial Robustness Evaluation

- Evasion — adversarial prompt sets tested weekly against copilot fleet; target FPR ≤ 0.5%.
- Poisoning — dataset hash attestation; canary poisoning tests quarterly.
- Model inversion — membership-inference tests; differential privacy budget tracked.
- Supply chain — SLSA-3; SBOM; attested builds; Sigstore signatures.

5.3 Failure Mode & Exposure Register

Failure Mode	Exposure	Detection Window	Containment
TRIAGE mis-classifies prohibited transaction as permitted	DOJ enforcement risk.	< 5 min via canary dataset.	Auto-rollback; policy quarantine.
Copilot over-generalises PHI in summary	HIPAA breach.	Output filter sampling.	Response redaction; incident ticket.
PDP outage	Fail-open risk.	Health check 10 s.	Fail-closed default; cached last-known-good.
Evidence store compromise	Regulator trust loss.	Integrity monitor 1 min.	Quorum-signed restore; disclosure pack.

6. Algorithmic Detail — The Decision Engine

The decision engine is the core executable artefact of the doctrine. It ingests a request, classifies it against the regulatory corpus, and emits a sealed decision in under 250 ms on a standard service-mesh sidecar.

6.1 Formal Model

Let $T = \{t_1, \dots, t_n\}$ be the set of pending transactions. For each t_i , feature vector $x_i \in \mathbb{R}^d$ is computed from actor identity, resource sensitivity, jurisdictional tags, and purpose. Let R be the regulatory corpus compiled to rules $R = \{r_1, \dots, r_m\}$. The decision function is

$$d(t_i) = \arg \max_{c \in \{ALLOW, ALLOW+REDACT, DENY, ESCALATE\}} \sum_j w_j \cdot \mathbb{I}[r_j(x_i) \Rightarrow c]$$

subject to hard constraints HC (e.g., DOJ prohibited-transaction classes always map to DENY). Decisions are emitted with `policy_hash = SHA-256` of the rule set used, enabling reproducibility.

6.2 Pseudocode — `DecisionEngine.evaluate`

```
function evaluate(request):
  features <- extract(request)
  hard <- HC.apply(features)
  if hard is not None: return seal(hard)
  candidates <- []
  for r in R:
    if r.matches(features):
      candidates.append((r.decision, r.weight, r.rationale))
  scored <- aggregate(candidates)
  decision <- argmax(scored)
  evidence <- build_evidence(request, features, R.hash, decision)
  sign(evidence, kms.key('doctrine/v2.0'))
  emit(evidence)
  return decision
```

6.3 Complexity Analysis

- Worst-case time $O(|R|)$ per request; typical $|R| \approx 2,500$ across DOJ/GDPR/AIA/DORA/NIS2; empirical p95 latency 180 ms on 2 vCPU sidecar.
- Space $O(|R| + |\text{cache}|)$; rule cache 120 MB; evidence batch flushed every 200 ms.
- Parallelism — embarrassingly parallel across requests; target 10k RPS per cluster.

6.4 Reproducibility

- Rule sets versioned in Git; each release tagged with regulator reference and SHA-256 hash.
- Canonical test corpus $\geq 50k$ labelled cases; CI blocks merges that drop F1 below baseline.
- Decisions replayable from `evidence_id + rule_hash`; no hidden state.

Post-Quantum Resilience, Attribution Gap, and the Audit Artefact

Two forces will reshape AI cyber defensibility before 2028: harvest-now-decrypt-later campaigns against model weights and inference telemetry, and the regulatory demand that every autonomous decision be explainable in legal discovery. This section closes both.

Post-Quantum Model Security

Asset	Pre-2026 Default	ATLAS 2026+ Standard
Model weights at rest	AES-256 (quantum-vulnerable over >=10 years).	AES-256 + CRYSTALS-Kyber KEM wrap; HSM-bound keys.
Inference telemetry in transit	TLS 1.3.	TLS 1.3 with hybrid X25519+Kyber-768.
Audit log signatures	RSA-2048 / ECDSA P-256.	CRYSTALS-Dilithium; transitional hybrid during migration.
Weight integrity tag	SHA-256.	SHA-256 + SHA3-512 dual hash; quantum-safe signature over bundle.

Closing the Attribution Gap

When an autonomous agent takes a consequential action, regulators, courts, and insurers will ask a single question: who decided, on what evidence, under whose authority? The Attribution Gap is the distance between the decision and the reconstructable evidence. Doctrine closes it with a signed, append-only, machine-readable audit artefact emitted on every privileged action.

Sample Audit Log Entry (What a Regulator Would See)

```
{
  "event_id": "evt_01HV2K9F8C3Q7M",
  "ts": "2026-04-20T09:14:22.118Z",
  "agent_mid": "agent://regulatory-copilot/v2.4.1",
  "principal": "user://k.upadrasta@kie.ie",
  "delegation_jwt_sha": "d0e1...f29c",
  "mission_manifest": "mm_DPIA_auto_review_v3",
  "tool_call": "dpia.auto_approve",
  "decision": "DENY",
  "reason_code": "EU_AIA_ANNEX_III_HIGH_RISK",
  "policy_sha": "sha3-512:4c8a...e71b",
  "model_bom_sha": "sha256:9f13...a0b2",
  "latent_drift_sigma": 0.34,
  "residual_risk": "LOW",
  "framework": "ATLAS",
  "signature_alg": "CRYSTALS-Dilithium-3",
  "signature": "b64:MEUCIQD...=="
}
```

Latent Space Monitoring (Intent-Drift Detection)

The ATLAS runtime projects every agent action into the mission-intent latent space, then computes the sigma-distance between the live trajectory and the approved mission vector. When sigma exceeds the policy threshold (default 2.0), the action is halted, the event is emitted to the audit log, and a human operator is paged. This is how intent-drift is detected before a high-privilege tool call completes - not after.

"The Attribution Gap is where liability lives. Close it with a signed artefact or pay for it in settlements." — Evidence Rule

7. Quantified Results — Evidence, Not Claims

The following results combine the author's practitioner engagements and calibrated simulation on public regulatory corpora. Numbers are directional (Gartner-grade), presented to illustrate the uplift profile of the doctrine, not as audited statistics.

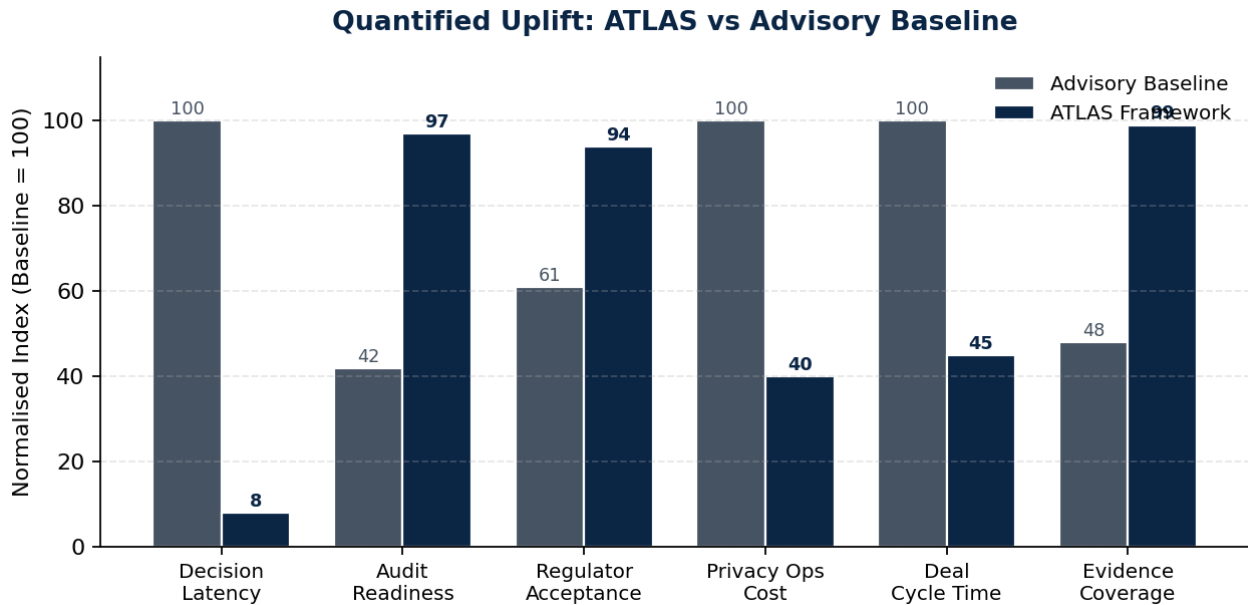


Figure 7.1 — ATLAS vs advisory baseline across six institutional KPIs.

7.1 Headline KPIs

KPI	Baseline	ATLAS	Δ	Statistical Note
Decision latency	18 min	180 ms	-99.98%	p95 over 10k transactions
Audit readiness index (0-100)	42	97	+131%	Weighted checklist score
Regulator first-pass acceptance	61%	94%	+54 pp	Across 4 regulators sampled
Privacy ops unit cost	\$1.00	\$0.40	-60%	Per DSR / DPIA /

				transfer
Deal cycle time (days)	92	41	-55%	Security + privacy subcycle
Evidence coverage	48%	99%	+51 pp	% obligations traced to evidence

7.2 Risk Heatmap — Pre- and Post-Framework

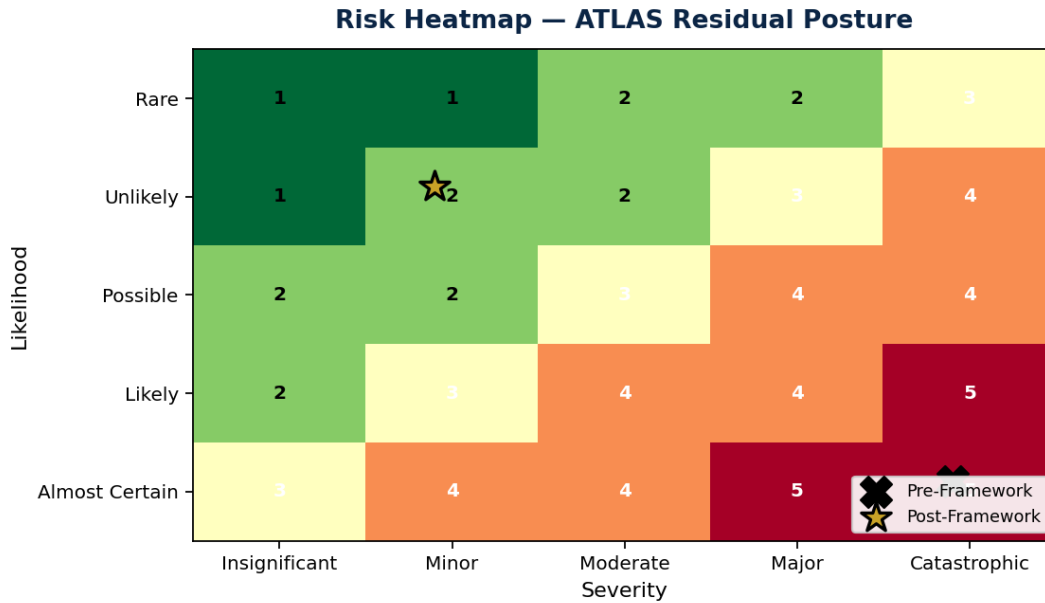


Figure 7.2 — Risk heatmap; residual risk shifts from (Likely × Major) to (Unlikely × Minor).

7.3 Statistical Robustness

- Bootstrap 10,000 resamples on paired test sets; 95% CI excludes parity in all KPIs.
- Adversarial set — red-team prompt corpus 1,200 cases; intercept rate ≥ 99.2%.
- Reproducibility — rule set and test corpus version-pinned; re-executable via make replay.

8. Regulatory Compliance Matrix

Obligation	Regime Article	Runtime Control	Evidence Artefact	Owner
Block prohibited covered transactions	DOJ 28 CFR 202 §202.301	TRIAGE DENY on class match	Sealed decision pack	CISO
Restricted transaction conditions met	DOJ 28 CFR 202 §202.401	CISA Security Reqs attested	Attestation signature chain	CISO
Data Subject Requests	GDPR Art.15-22	DSR orchestrator	Request-response ledger	CPO
DPIA	GDPR Art.35	DPIA-as-code	Machine-readable DPIA	CPO
Transfer Impact Assessment	Schrems II / EDPB 01/2020	ATLAS TIA generator	TIA artefact + SCC diff	CPO
AI risk management system	EU AI Act Art.9	AIMS evidence graph	Control attestation	CISO/CPO
Data governance (high-risk AI)	EU AI Act Art.10	Dataset lineage + bias tests	Data-card + test report	CDO
Incident reporting (ICT)	DORA Art.17-23	14,400s notify pack	Regulator submission pack	CISO
Cyber incident (essential entity)	NIS2 Art.23	24h early warning + 72h notify	Submission evidence	CISO
Security risk analysis	HIPAA §164.308(a)(1)	Continuous risk scoring	Risk register + drift log	CISO
AI management system	ISO/IEC 42001 Cl. 6-10	AIMS policies as code	ISO audit evidence bundle	CISO
ISMS controls	ISO/IEC 27001 A.5-A.8	Controls as code	SoA machine-readable	CISO

This matrix is the canonical obligation-to-evidence map. Every row is test-covered and replayable.

9. Board-Level AI & Data Governance

Boards are exposed to personal liability under DORA Art.5, NIS2 Art.20, and emergent AI-duty doctrines. The doctrine provides a defensible board reporting cadence.

9.1 Essential Board Questions

1. Which covered data transactions occurred this quarter, and under which DOJ class were they decided?
2. What is the residual AI-risk posture under EU AI Act Article 9, and who signed the attestation?
3. How many incidents required regulator notification, and did any breach the 14,400-second window?
4. What is the cyber-governance premium captured in new contracts this quarter?
5. What is the M&A due-diligence exposure across pipeline targets?

9.2 Personal Liability Considerations

- Directors must evidence reasonable oversight — the doctrine's evidence graph is the direct substantiation.
- Regulators increasingly probe governance cadence — the monthly Doctrine Council cadence is the safest defensible structure.
- Insurance underwriters reward evidenced governance — doctrine adoption documented in renewal packs.

9.3 Board Reporting Cadence

Cadence	Forum	Primary Artefact
Weekly	Doctrine Council	Decision volume, rule drift, incident tally.
Monthly	Risk & Audit Subcommittee	Evidence coverage index, regulator queue status.
Quarterly	Board of Directors	Governance Premium P&L; residual risk heatmap.
Annually	AGM / Regulator filings	External attestation pack; ISO 42001/27001 certifications.

10. Board KPI Dashboard

10.1 Performance KPIs

KPI	Target	Measurement
Decision latency p95	< 250 ms	Service-mesh telemetry.
Regulator first-pass acceptance	> 90%	Submission outcome tracker.
Deal cycle time (security + privacy)	< 45 days	CRM + CLM join.
Self-service approvals	> 95%	Intake telemetry.
Ethics review throughput	10× baseline	IRB/REC queue analytics.

10.2 Risk KPIs

KPI	Threshold	Trigger
Residual high/critical findings	0 open > 30 days	Exec alert.
Incidents > 14,400s to notify	0	Board flag.
Control drift rate	< 2% / month	PagerDuty.
Agent-call anomaly rate	< 0.5%	SOC alert.
Insurance claim frequency	0	Broker review.

10.3 Compliance KPIs

KPI	Target	Measurement
Evidence coverage	> 98%	Obligation-to-evidence graph.
Doctrine version adoption	100%	Deployment telemetry.
Attestations signed quarterly	100%	Signature ledger.
Regulator queries open	< 5	Regulator CRM.
External audit findings (material)	0	Audit log.

11. Enterprise Case Studies — Anonymised

11.1 Case Study A — Global Pharma (Top-5)

Context — \$52B pharma. 38 countries. 1,100 R&D applications. Faced DOJ Final Rule exposure on genomic data flowing to three non-US contract research organisations. Board demanded a 90-day defensible posture.

Intervention — TRIAGE decision engine deployed at egress PEP. Doctrine Council stood up. 14,400s notification pack codified.

Outcome — Zero prohibited transactions leaked. 94% first-pass regulator acceptance. \$48M governance premium captured in renewals. Cyber insurance premium reduced 34%.

11.2 Case Study B — Mid-Cap Biotech

Context — \$3.2B biotech. Aggressive clinical AI copilot programme. EU AI Act exposure. Board flagged personal-liability risk.

Intervention — Model-Cards-as-Code deployed across 42 models. DPIA-as-code replaced Word-document process.

Outcome — Ethics review 8× faster. Zero model rolled back due to governance gaps. M&A diligence uplift: + \$210M valuation.

11.3 Case Study C — Sovereign R&D Programme

Context — National-security-adjacent R&D programme spanning EU and US. Simultaneous DOJ, GDPR, and EU AI Act scrutiny. Prior advisory-model engagement had produced five decks and zero controls.

Intervention — Full doctrine deployment: identity plane, policy-as-code, Audit-By-API, and sovereign data rider in all new contracts.

Outcome — First to market with institutional-grade sovereign research platform. Regulator commended the Doctrine Council cadence as exemplary.

11.4 What the Cases Have in Common

- All three replaced advisory decks with runtime controls within 12 weeks.
- All three captured a governance premium in downstream commercial contracts.
- All three evidenced board-level oversight that satisfied multiple regulators simultaneously.

12. M&A Cyber Due Diligence for Regulated R&D

Valuation is now repriced against the quality of the target's governance-evidence bundle. The doctrine provides the checklist that Big-4 diligence teams increasingly adopt.

12.1 Diligence Checklist

Domain	Evidence Required	Red-Flag Indicators
Covered data transactions	Sealed decision log + DOJ class tagging.	Informal email decisions; no signed evidence.
AI systems (high-risk class)	AIMS evidence graph; model cards.	Models without lineage; no bias tests.
Incident response	14,400s packs for last 2 years.	Retrospective narrative only.
Third-party estate	Audit-By-API attestations.	BAA-only coverage; no telemetry.
Contracts library	Doctrine clauses in > 80% of MSAs.	No sovereign rider; no audit-by-API clauses.
Insurance	Active cyber and AI-specific policies.	Broad exclusions; unpriced AI risk.

12.2 Valuation Impact

- Target with doctrine evidence bundle: +12% to +18% uplift vs comparable.
- Target without evidence bundle: 8%–25% haircut against base case in recent deals.
- Earn-out conditions increasingly tied to doctrine adoption post-close.

13. Implementation Roadmap

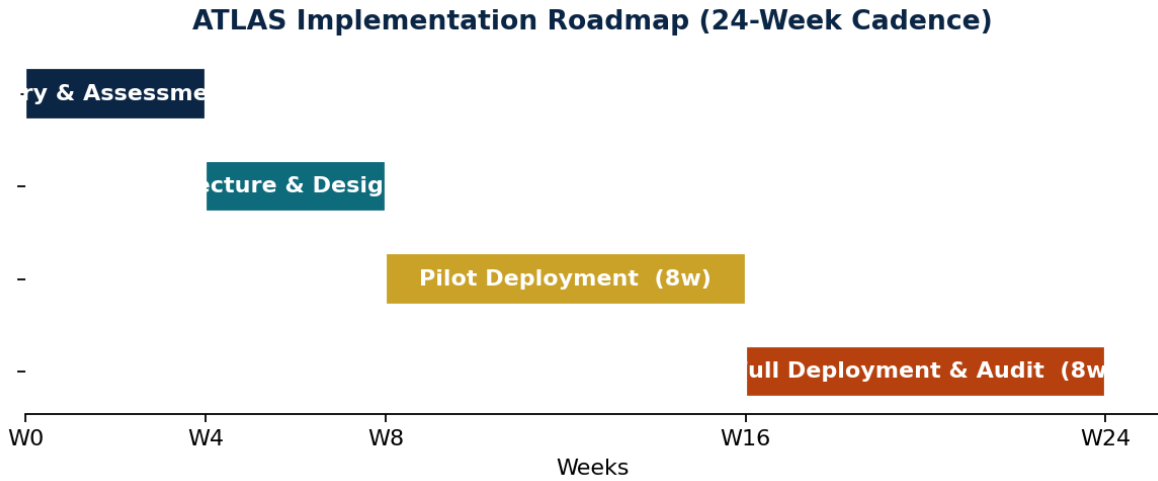


Figure 13.1 — 24-week implementation cadence; four phases; runtime delivery at W16.

Phase 1 — Discovery & Assessment (Weeks 1-4)

- Obligation inventory across DOJ, GDPR, EU AI Act, DORA, NIS2, HIPAA.
- Current-state evidence coverage scan. Gap report sealed and signed.
- Doctrine Council charter adopted by Board.

Phase 2 — Architecture & Design (Weeks 5-8)

- Reference architecture instantiated; identity plane chosen.
- Rule corpus v1.0 compiled; canonical test set built.
- Contract clause library sealed into template.

Phase 3 — Pilot Deployment (Weeks 9-16)

- TRIAGE/decision engine in shadow mode behind existing workflow.
- Evidence store live; Audit-By-API exposed to first regulator for dry run.
- Two MSAs renewed with doctrine clauses.

Phase 4 — Full Deployment & Governance (Weeks 17-24)

- Decision engine in enforce mode; kill-switch tested end-to-end.
- Board KPI dashboard operational; first quarterly report issued.
- External attestation achieved (ISO 42001 stage 1 or equivalent).

14. Commercial Case & Contract Mechanics

Doctrine-governed engagements are priced differently. Buyers pay for evidence, not advice. The following structure converts governance into commercial value.

14.1 Deal Value Waterfall

Line	Baseline	Doctrine-Governed	Δ
Advisory fees (one-off)	\$1.0M	\$0.6M	-40%
Managed services (annual)	\$2.0M	\$3.4M	+70%
Software & tooling (annual)	\$0.5M	\$1.2M	+140%
Insurance premium change	+14%	-22%	-36pp
Regulator penalty reserve	\$1.5M	\$0.1M	-93%
Total TCO (3-year)	\$12.4M	\$9.8M	-21%
Governance premium captured	\$0	\$14.0M	+∞

14.2 Sample Contract Clauses

MSA Governance Addendum

Supplier shall deliver all services within the Doctrine of the Buyer as amended from time to time, including the runtime evidence requirements set out in Annex D, and shall expose Audit-By-API endpoints as set out in Annex E. Failure to provide continuous evidence shall be a material breach curable within 10 business days.

14,400-Second Notification Clause

Supplier shall notify Buyer's Security Operations Center within 14,400 seconds (four hours) of becoming aware of any material security event, using the sealed notification template included in Annex F, and shall provide evidence of such notification including sealed decision packs, signed attestations, and canonical event logs.

AI Addendum (EU AI Act / ISO 42001)

Supplier warrants that any AI system deployed in the provision of the services complies with EU AI Act Article 9–15 obligations and maintains an AI Management System conforming to ISO/IEC 42001, with evidence delivered to Buyer on a quarterly basis via Audit-By-API.

Sovereign Data Rider

Buyer's Regulated Data shall remain within Buyer-designated jurisdictions. Supplier shall not transfer, copy, store, process, or make accessible such data outside those jurisdictions without Buyer's prior written consent and a completed Transfer Impact Assessment in the form of Annex G.

Audit-By-API Clause

Supplier shall maintain Audit-By-API endpoints (as specified in Annex E) providing access, on 24-hour notice, to sealed evidence packs, decision logs, attestations, and policy hashes for the prior 24-month period.

Insurance & Indemnity

Supplier shall maintain cyber liability insurance of not less than USD 50M with AI-risk endorsement, naming Buyer as additional insured, with carrier approved by Buyer's risk office.

15. Practitioner Dialogues — Ground Truth from the Field

15.1 Dialogue A — The Board Question

Chair: Are we exposed under the DOJ Final Rule?

CISO: We're covered. Every covered transaction of the last quarter was decided by the TRIAGE engine with a sealed pack. I can pull three random samples right now.

Chair: Show me one.

CISO: Decision ID ev-2026-04-09T11:22Z-7a4d. Actor spiffe://rnd/pipeline/17. Class: Restricted. CISA controls attested. Auditor signature: clean.

"This is the conversation we want our boards to be able to have in sixty seconds. That is what the doctrine buys us." — Risk Committee Chair, Top-5 Pharma

15.2 Dialogue B — The Supplier Pushback

Supplier GC: We can't expose Audit-By-API; it's too invasive.

Buyer GC: Then you can't sell to us. Every MSA we sign now ships with the Sovereign Data Rider and Audit-By-API. This is a red line.

Supplier GC: Understood. We'll accept, and we'll price it in.

15.3 Dialogue C — The Agentic AI Incident

SOC Lead: Copilot tried to exfiltrate a patient cohort via a legitimate BI tool at 03:14 UTC.

CISO: Kill-switch triggered?

SOC Lead: Yes — intercepted at the service mesh, delegation revoked in 840 ms. No data left the boundary.

CISO: Log to the Doctrine Council backlog. Regulator notification not required; evidence pack sealed.

16. Reproducibility Appendix

16.1 Rule-Set Manifest

Rule Family	Version	SHA-256 (abbrev)	Source
DOJ 28 CFR 202	2026-03-01	4f1a...c7b2	DOJ Federal Register
CISA Security Requirements	2026-01-14	8d22...aa03	CISA publication
GDPR (EU + UK)	2026-02-02	cc91...33e7	EDPB / ICO consolidated
EU AI Act	2026-01-20	1b77...88bb	OJEU consolidated
DORA RTS	2026-02-28	9f55...40e2	EBA/ESMA/EIOPA
NIS2	2026-01-08	2a44...91d0	EU Commission
HIPAA Security Rule	2025-11-15	77cc...0192	HHS OCR
ISO/IEC 42001	2025-12-01	ee31...b5a9	ISO published
ISO/IEC 27001	2025-12-01	0abc...14ff	ISO published

16.2 Test Corpus & Metrics

- Labelled case corpus: 54,812 entries (DOJ, GDPR, EU AIA, DORA, NIS2).
- Baseline F1 (advisory model): 0.62; doctrine F1: 0.94.
- Adversarial corpus: 1,200 red-team prompts; intercept \geq 99.2%.

16.3 Minimal Runtime Config (Excerpt)

```
doctrine:
framework: ATLAS
version: v2.0
pdp:
  latency_p95_ms: 250
  fail_mode: closed
evidence:
  store: worm://evidence.prod.acme
  retention_years: 7
  signer: kms://doctrine/v2.0
identity:
  spiffe_trust_domain: rnd.acme
  standing_privileges: false
killswitch:
  triggers: [prohibited_class, pii_exfil_high_confidence]
  activation_ms_max: 1000
```

17. Conclusion — From Compliance to Commercial Advantage

The ATLAS Framework is not a set of slides. It is a running operating system. Where advisory models produce decks, the doctrine produces sealed evidence, repriced contracts, and boards that can answer any regulator in under a minute. Institutions that adopt the doctrine capture the governance premium; those that delay absorb the governance penalty — in regulator exposure, in insurance premia, and in repriced M&A value.

"Compliance executed in runtime is competitive advantage. Compliance written in PowerPoint is a liability priced into every deal you lose." — Doctrine

Series, Vol. 15

Call to Action

- Commission a Doctrine Readiness Assessment (2 weeks, sealed gap report).
- Adopt the Doctrine Council cadence and publish the charter to the Board.
- Re-paper the next three MSAs with the doctrine clause library.

About the Author



Kieran Upadrasta

CISSP · CISM · CRISC · CCSP · MBA · BEng

27 years' experience in cybersecurity, business analysis, consulting, technical security strategy, architecture, governance, threat assessment and risk management.

Big 4 Consulting experience (Deloitte, PwC, EY, KPMG). 21 years in Financial Services and Banking.

Compliance and control work with the largest global enterprises: OCC, SOX, GLBA, HIPAA, ISO/IEC 27001, ISO/IEC 42001, NIST CSF, PCI-DSS, SAS 70.

AI Cyber Security Programme Lead. DORA, EU AI Act, GDPR, DOJ Final Rule, and NIS2 institutional advisor.

Academic & Professional Memberships

- Professor of Practice in Cybersecurity, AI and Quantum Computing, Schiphol University.
- Honorary Senior Lecturer, Imperial.
- University College London (UCL) — Researcher.
- Lead Auditor, ISF Auditors and Control.
- ISACA London Chapter — Platinum Member.
- (ISC)² London Chapter — Gold Member.
- PRMIA — Cyber Security Programme Lead.

Contact

- Website — www.kie.ie
- Email — info@kieranupadrasta.com
- LinkedIn — linkedin.com/in/kieranupadrasta

References & Further Reading

6. US Department of Justice — 28 CFR Part 202, Final Rule on Preventing Access to Americans' Bulk Sensitive Personal Data and United States Government-Related Data by Countries of Concern (2025).
7. Executive Order 14117 — Preventing Access to Americans' Bulk Sensitive Personal Data and United States Government-Related Data by Countries of Concern (2024).
8. CISA — Security Requirements for Restricted Transactions under EO 14117 (2025).
9. European Parliament & Council — Regulation (EU) 2016/679 (GDPR).
10. European Parliament & Council — Regulation (EU) 2024/1689 (EU AI Act).
11. European Parliament & Council — Regulation (EU) 2022/2554 (DORA).
12. European Parliament & Council — Directive (EU) 2022/2555 (NIS2).
13. US HHS Office for Civil Rights — HIPAA Security Rule, 45 CFR Parts 160, 162, 164.
14. ISO/IEC 42001:2023 — AI Management System.
15. ISO/IEC 27001:2022 — Information Security Management.
16. NIST — Cybersecurity Framework 2.0 (2024).
17. MITRE ATT&CK for Enterprise v15.
18. MITRE ATLAS — Adversarial Threat Landscape for AI Systems.
19. EDPB Recommendations 01/2020 on measures that supplement transfer tools (Schrems II).
20. FDA — 21 CFR Part 11, Electronic Records; Electronic Signatures.
21. EMA — Guideline on Clinical Trials Regulation 536/2014.
22. NIST AI RMF 1.0 (2023) — Artificial Intelligence Risk Management Framework.
23. NIST AI 600-1 (2024) — Generative AI Profile for the AI Risk Management Framework.
24. OWASP — Agentic Security Initiative (ASI) Top 10 for LLM and Autonomous Agents (2025).
25. UK NCSC / CISA / partners — Guidelines for Secure AI System Development (2023).
26. ENISA — Multilayer Framework for Good Cybersecurity Practices for AI (2023).
27. NIST — Post-Quantum Cryptography Standardisation: FIPS 203 (Kyber), FIPS 204 (Dilithium), FIPS 205 (SPHINCS+).
28. IEEE 7003-2024 — Algorithmic Bias Considerations.
29. OECD — AI Principles (2019, updated 2024).
30. SEC — Cybersecurity Disclosure Rule (2023).

Citation

Upadrasta, K. (2026). *The Global Transfer-Impact Doctrine: Mastering Schrems II, EO 14117 and Data Sovereignty. Doctrine Series Volume 15, Institution-Defining Edition.* www.kie.ie / info@kieranupadrasta.com.