THE $3 BILLION WARNING

Governance Failures in Automated Finance and the Deterministic Infrastructure Response

Tanishq Dasari
Independent Researcher · AnimusLab
tan@animuslab.dev · github.com/AnimusLab/anchor · animuslab.dev

$3B+	5	5	3.8+
Total Capital at Risk (5 Cases)	Regulatory Frameworks Satisfied	Months to Reconstruct Flash Crash	Python Compatibility

Abstract

Five documented failures in automated financial systems have collectively destroyed or misallocated more than $3 billion in institutional capital. In each case, the technology that caused the failure existed and was understood. The infrastructure to govern it in real time did not. This paper examines five case studies - Knight Capital Group (2012), the Flash Crash (2010), Zillow Offers (2021), TSB Bank (2018), and Citibank's Revlon transfer (2020) - and extracts a common structural pattern: these failures were not caused by systems doing unexpected things. They were caused by systems doing exactly what they were configured to do, in the absence of governance infrastructure capable of enforcing what they were authorized to do. We examine the current regulatory enforcement landscape - the EU AI Act (enforcement begins August 2026), CFPB Regulation B and Circular 2024-03 on AI-assisted adverse action notices, and SEC 2026 examination priorities - and demonstrate that five major frameworks converge on one technical requirement: a tamper-evident, per-decision record of what an automated system decided, on what basis, under what authority, and why. We present Anchor, an open-source governance engine satisfying this requirement across all five frameworks simultaneously - a property we term regulatory polyglottism.

Keywords: AI governance, execution governance, financial services, regulatory compliance, EU AI Act, CFPB Regulation B, audit trails, deterministic enforcement, agentic AI

JEL Classification: G18, G20, K22, K24, O33

Availability: pip install anchor-audit · github.com/AnimusLab/anchor · anchorgovernance.tech

Revision note (this version). This revision corrects a citation error in Sections 3.2 and 4.4 regarding the CFPB's October 2024 enforcement action against Goldman Sachs and Apple, and replaces the performance benchmarks in Section 4.6 with independently reproducible benchmark runs against anchor-audit, executed via the project's own benchmarks/bench_anchor.py harness on two distinct hardware/OS environments - each run more than once, with the observed range reported rather than a single figure. All tables in this revision use native Word table structures rather than fixed layout, so they retain formatting when edited. No other section has been altered. Prior version: distributed 2026-06-13.

1. Introduction: The Pattern Behind the Losses

Between 2010 and 2021, five automated financial systems caused or contributed to losses and regulatory costs exceeding $3 billion in aggregate. The systems involved were not experimental. They were production-grade infrastructure operated by sophisticated institutions - a market maker ranked among the largest in U.S. equities, a publicly traded real estate platform with billions in capital, a global systemically important bank, and a major UK retail bank undergoing a supervised IT migration.

In each case, the failure was investigated, documented, and attributed. In each case, the attribution pointed not to model inaccuracy or traditional software bugs, but to a common structural absence: the institution could not produce, in real time or after the fact, a complete and verifiable record of what its automated systems were authorized to do, what they actually did, and whether the two matched.

This paper argues that this absence - which we term the execution governance gap - is distinct from model governance (does the model produce accurate outputs), data governance (is data classified correctly), and content governance (is the output appropriate). Execution governance asks a more fundamental question: was this action authorized, and can we prove it?

The Core Distinction. “Is this output safe?” is a question about model quality. It is probabilistic and genuinely difficult. “Was this action authorized, and can we prove it?” is a question about governance and provenance. It is answerable today, with existing engineering techniques, without requiring any advance in the underlying models.

This question is answerable today with existing engineering techniques - cryptographic hashing, abstract syntax tree analysis, append-only audit logs, and structural output validation. The five cases examined below each illustrate a failure that execution governance infrastructure would have detected, blocked, or made immediately reconstructable.

1.1 The Distinction That Matters

Current AI governance tooling addresses important questions. Model monitoring platforms ask whether output quality is degrading. Bias classifiers ask whether decisions are discriminatory in aggregate. Content filters ask whether a specific output is appropriate. Each is valuable.

None of them asks what regulators are actually enforcing: for this specific decision, at this specific moment, was the action authorized under the institution's stated policy, and does a verifiable record of

that authorization exist? This is the question Anchor is built to answer - a question about governance and provenance, not model quality.

2. Five Failures and the Governance Gap

The five cases below were selected because each is publicly documented at a level of detail supporting rigorous counterfactual analysis, and because each represents a distinct failure mode within the execution governance taxonomy.

Case	Year	Proximate Cause	Governance Gap	Total Cost
Knight Capital	2012	Deployment inconsistency + dormant code	No deprecated code detection; no fleet integrity verification	$460M+ losses; $12M SEC fine
Flash Crash	2010	Algorithmic feedback loop	No per-decision audit trail; 5-month reconstruction	Multi-participant systemic disruption
Zillow Offers	2021	Model drift under market change	No per-decision reasoning record; no exposure escalation	$500M+ write-downs
TSB Bank	2018	Incomplete deployment verification	No configuration drift detection; no go/no-go attribution	£600M+ total impact
Citibank Revlon	2020	Authorization gap in payment system	No structured authorization record; no amount-range enforcement	$500M permanently lost

2.1 Knight Capital Group - August 1, 2012

On August 1, 2012, Knight Capital Group - at the time one of the largest market makers in U.S. equities, executing approximately 10–15% of daily trading volume in NYSE and Nasdaq-listed stocks - deployed new software to support the NYSE's Retail Liquidity Program across eight production trading servers. The deployment required updating code on all eight servers simultaneously. It reached seven of them.

The eighth server, through a single technician's oversight and the absence of any second-person verification requirement, continued running the prior version of the codebase, which contained a dormant function called Power Peg - internal trading logic Knight had ceased using in 2003. For 45 minutes, Power Peg sent continuous streams of child orders across 154 stocks without limit. By the time Knight's direct market access was terminated, the firm had executed over 4 million trades, accumulating approximately $3.5 billion in long positions and $3.15 billion in short positions. Net loss: more than $460 million.

The SEC's administrative order (Release No. 70694, October 2013) found violations of Exchange Act Rule 15c3-5 and Regulation SHO, imposing a $12 million civil penalty, and noted that a written protocol requiring retesting of Power Peg when it was last modified in 2005 could have identified that its position safeguard had been disabled.

What Execution Governance Changes. Any one of three mechanisms interrupts this failure path before a single live order is placed: static constitutional analysis producing a BLOCKER finding for any deployment artifact containing a blocklisted deprecated module; cryptographic fleet integrity verification preventing

any server running unsealed code from entering production; and behavioral verification confirming whether a code path behaves consistently with its specification under execution.

2.2 The Flash Crash - May 6, 2010

On May 6, 2010, U.S. equity markets fell approximately 9% within minutes before recovering most of the loss in the same session. The joint SEC-CFTC staff report required five months to produce a complete account of an event that lasted approximately 36 minutes in its most severe phase. A separate actor, Navinder Singh Sarao, was later found to have engaged in systematic spoofing that amplified the volatility, and was not charged until nearly five years after the event.

> *The Core Governance Gap. No participant's algorithmic trading systems produced a structured, per-decision record at the time decisions were made. The claim here is narrow: Anchor's runtime Decision Audit Chain does not claim it would have shortened the market-wide five-month investigation. The claim is that for any future investigation concerning a specific institution's own systems, structured per-decision records transform forensic reconstruction from a multi-week manual process into a database query.*

2.3 Zillow Offers - 2021

Zillow Offers was an algorithmic home-buying program whose pricing models continued generating purchase offers after the historical price relationships they relied on stopped holding under 2021's pandemic-driven housing volatility. By Q3 2021 Zillow disclosed write-downs exceeding $300 million in a single quarter, ultimately over $500 million across the wind-down, and shut the program down entirely, laying off approximately 25% of its workforce.

This was, at root, a model accuracy failure, and execution governance does not solve model accuracy. What was also absent was a continuously queryable per-decision record of the model's stated reasoning, and a constitutional requirement that purchases above a defined exposure threshold trigger structured human review.

2.4 TSB Bank IT Migration - April 2018

TSB Bank's 2018 migration off Lloyds-inherited infrastructure onto a Banco Sabadell-built platform locked a significant proportion of TSB's 5.2 million customers out of their accounts, in some cases for months. An independent review found approximately 2,000 defects at go-live; the board had been informed of approximately 800. Only one of two data centers had been tested before launch. The FCA and PRA jointly fined TSB £48.65 million in December 2022; combined with customer redress and operational costs, total impact exceeded £600 million.

What Execution Governance Changes. Static constitutional analysis would surface the untested dual-datacenter configuration as a BLOCKER before go-live. A structured readiness record requirement would have surfaced the 2,000/800 discrepancy at the decision point. The governance floor invariant ensures any decision to proceed despite open findings is individually attributed and sealed at the moment it is made, rather than four years later under regulatory investigation.

2.5 Citibank Revlon Transfer - August 11, 2020

A Citibank employee configuring a routine $7.8 million interest payment on an $894 million syndicated Revlon loan selected an option that caused the system to transmit the full principal balance instead. Three reviewers confirmed the payment without catching the error, because the authorization system presented them with a payment amount, not a structured record of what they were authorizing. Approximately $500 million was never recovered after a federal court ruled in the lenders' favor under the discharge-for-value doctrine.

What Execution Governance Changes. Structured Mode applied to any payment instruction above a defined threshold would require a machine-readable authorization record specifying payment type and amount relative to the authorized range. A payment of $893 million against an authorized cycle payment of $7.8 million - a ratio of 114:1 - produces a BLOCKER finding before transmission.

3. The Regulatory Enforcement Landscape

The five failures above produced regulatory responses across four jurisdictions. Those responses, read together, describe a convergent technical requirement.

3.1 EU AI Act (Regulation 2024/1689)

The EU AI Act classifies credit scoring, anti-money-laundering monitoring, and fraud detection as high-risk AI systems under Annex III. Article 12 requires automatic, tamper-evident logging throughout a system's lifetime. Enforcement on high-risk systems begins August 2026, with penalties reaching €30 million or 6% of global annual revenue, whichever is higher.

3.2 CFPB Regulation B and AI-Assisted Adverse Action

[Corrected in this revision - see Revision Note, p.1]

Regulation B (ECOA) requires that applicants denied credit receive a statement of specific, accurate reasons for the denial. In September 2023 the CFPB issued Circular 2024-03, extending these requirements explicitly to AI-assisted lending decisions: a creditor cannot satisfy adverse action notice requirements with reasons that are overly broad, that describe generic model categories rather than the specific factors that drove the decision, or that a consumer could not act on. A credit decision system returning unstructured prose rather than a specific, machine-readable reason code is, under this guidance, non-compliant by construction - independent of whether the underlying model is accurate or unbiased.

The CFPB has separately signaled it will treat opaque, “black box” credit models as a supervisory concern in their own right: then-Director Rohit Chopra stated in 2023 that federal anti-discrimination and adverse-action law applies in full to algorithmic underwriting, and that creditors may not use model complexity as a defense for failing to provide specific reasons for a denial. This is the correct regulatory basis for the individual-decision explainability requirement this paper argues Anchor's Structured Mode satisfies.

A note on the Goldman Sachs / Apple Card precedent. An earlier version of this paper cited the CFPB's October 23, 2024 enforcement action against Goldman Sachs and Apple ($89 million combined penalty and redress) as an AI credit-decision explainability precedent. On review, this was incorrect. That order concerned Apple Card transaction-dispute handling: Apple's failure to forward tens of thousands of consumer disputes to Goldman, and Goldman's failure to investigate the disputes it did receive, in violation of the Consumer Financial Protection Act and the Truth in Lending Act - a servicing and dispute-resolution failure, not an algorithmic credit-decision explainability finding. It is a legitimate data point for this paper's broader argument that regulators are actively enforcing against process and structure failures at Goldman specifically, and we retain it below for that purpose - but it is not evidence for the specific No-Prose Rule claim made in Section 4.4, which now rests on Regulation B and Circular 2024-03 directly.

Separately, in 2019 Apple Card drew public scrutiny after individual users, including Apple co-founder Steve Wozniak, reported that the algorithm underlying credit-limit decisions appeared to offer women lower limits than men with similar financial profiles despite shared accounts. That controversy prompted an investigation by the New York State Department of Financial Services, which concluded in 2021 that it had found no evidence of intentional or disparate-impact discrimination under New York law, while noting that the algorithm's inputs and logic were not independently reproducible by regulators from Goldman's documentation alone. We cite this only as motivating context for why individual-decision explainability and reproducibility - not aggregate fairness metrics alone - is the standard both the public and regulators are increasingly applying to algorithmic credit decisions; it is not offered as a finding of a legal violation.

3.3 SEC 2026 Examination Priorities

The SEC's Division of Examinations named AI governance a top priority for 2026 examinations of registered investment advisers and broker-dealers, including the accuracy of AI marketing claims (“AI washing”) and audit trails of AI systems used in trading and investment decisions.

3.4 Reserve Bank of India FREE-AI Report (August 2025)

The RBI's Framework for Responsible and Ethical Enablement of AI issued 26 recommendations, including per-decision audit trails reportable to the CIMS portal (Rec. 7), individual-decision explainability with machine-readable reason codes (Rec. 14), and mandatory algorithmic fairness audits for credit and lending systems (Rec. 19).

3.5 NIST AI Risk Management Framework

The NIST AI RMF establishes Govern, Map, Measure, and Manage as core functions, requiring documented policies, accountability structures, and continuous behavioral monitoring - an ongoing operational requirement rather than a one-time assessment.

3.6 The Convergence

Read individually, these are five separate compliance obligations from five separate bodies across four jurisdictions. Read together, they describe a single underlying requirement in five regulatory dialects:

An institution must be able to produce, on demand, a record of what an AI system decided, on what basis, under what authority, and why - and that record must be tamper-evident.

4. Anchor: The Deterministic Infrastructure Response

4.1 Architecture Overview

Anchor is an open-source governance engine (pip install anchor-audit, Apache 2.0 license, Python 3.8+) enforcing a single, cryptographically signed rule set - a constitution - across two layers. Layer 1 (Static) parses AI-adjacent source code into abstract syntax trees using Tree-sitter and evaluates those trees against the constitution's rules in CI, as a pre-commit hook, or on demand. Layer 2 (Runtime) intercepts live AI inference via a lightweight decorator, hashes the input and output, evaluates the output against the constitution, and writes the result to an append-only, cryptographically chained audit log - the Decision Audit Chain (DAC).

4.2 The Decision Audit Chain

Each intercepted runtime call produces an AuditEntry: timestamp, model provider and version, active jurisdiction, git commit of the calling code, compliance status, SHA-256 hashes of the input and output, and specific rules evaluated. Entries form a hash-chained, HMAC-SHA256-signed, append-only log in which altering any historical entry invalidates every entry that follows.

4.3 Regulatory Polyglottism

The AuditEntry exposes jurisdiction-specific output methods - adverse_action_reasons() (CFPB Regulation B), cims_payload() (RBI CIMS format), eu_article12_record() (EU AI Act Article 12) - translating one underlying record into the format a given regulator requires. A single decision can be rendered simultaneously as a CFPB-compliant adverse action notice, an RBI CIMS submission, and an EU AI Act Article 12 log entry, because the frameworks are independently demanding the same underlying artifact in different regulatory language. We term this regulatory polyglottism.

4.4 The No-Prose Rule

[Corrected in this revision - see Revision Note, p.1]

Regulation B and CFPB Circular 2024-03 require that an AI-assisted adverse action notice state the specific factors that actually drove a denial, in terms a consumer can act on - not a generic description of the model's category or a narrative explanation that cannot be traced to specific input features. Anchor's Structured Mode enforces this structurally: in Structured Mode, a decision is not valid unless it includes a machine-readable ReasonCode and FeatureAttribution. Returning unstructured prose for a high-stakes decision is itself a BLOCKER finding by construction - the absence of the required structure is detected with certainty because it is a question about the shape of the output, not its content.

4.5 GOVERNANCE.lock and the Governance Floor Invariant

GOVERNANCE.lock seals the constitutional files via SHA-256 hash, verified on every scan; any modification produces a hash mismatch that aborts the scan under strict mode. The governance floor

invariant permits local policy overrides to raise a rule's severity above the constitutional floor, but rejects any attempt to lower it - enforced at the engine level, not as a policy convention.

4.6 Performance

[Replaced in this revision - see Revision Note, p.1] A governance layer that adds material latency to every model call is not viable at institutional scale. The figures below are produced by the project's own published benchmark harness, benchmarks/bench_anchor.py, run against anchor-audit's current main branch with 10,000 timed iterations per suite following a 200-iteration warmup, exactly as originally specified. Each environment below was run independently more than once, and we report the observed range rather than a single point estimate - both because that is a more honest representation of real-world latency, and because a single figure invites exactly the kind of scrutiny this revision exists to resolve. All runs are independently reproducible by cloning github.com/AnimusLab/Anchor and executing: python benchmarks/bench_anchor.py --iterations 10000 --json.

Environment A - Linux (Ubuntu 24.04.4 LTS, Python 3.12.3, Intel Xeon @2.80GHz, 1 shared vCPU - cloud/container instance; 3 independent runs)

Operation	Mean (range across runs)	P95	P99	Throughput
Policy Evaluation (118 rules, 713-byte fixture)	4.23 – 5.36 ms	4.52 – 5.80 ms	6.11 – 7.48 ms	187 – 236 / sec
Hash Generation (SHA-256, 33-byte payload)	~1.0 µs	1.4 – 1.6 µs	2.9 – 3.0 µs	~1,000,000 / sec
Receipt Generation (ETH scan + chain hash + JSONL write)	1.51 – 2.17 ms	1.65 – 2.37 ms	2.09 – 3.13 ms	461 – 663 / sec
Replay Verification (HMAC-SHA256)	4.5 – 5.3 µs	4.7 – 8.3 µs	11.1 – 13.0 µs	188,000 – 220,000 / sec

Environment B - Windows (Windows 11 Home, Python 3.14.4, Intel Core i5-12450H - consumer laptop; 2 independent runs under varying background system load)

Operation	Mean (range across runs)	P95	P99	Throughput
Policy Evaluation (118 rules, 713-byte fixture)	4.30 – 11.92 ms	4.79 – 17.46 ms	5.33 – 53.31 ms	84 – 232 / sec
Hash Generation (SHA-256, 33-byte payload)	0.75 – 2.98 µs	0.80 – 2.60 µs	0.90 – 8.60 µs	335,800 – 1,331,115 / sec
Receipt Generation (ETH scan + chain hash + JSONL write)	10.11 – 25.57 ms	12.58 – 34.75 ms	14.24 – 44.55 ms	39 – 99 / sec
Replay Verification (HMAC-SHA256)	3.40 – 8.24 µs	3.70 – 9.10 µs	4.00 – 30.40 µs	121,323 – 294,108 / sec

The two Windows runs differ by a consistent 2.4–4x factor across all four operations simultaneously, in the same direction - the signature of shared machine load (background processes, real-time antivirus scanning, or thermal/power state) rather than a defect in any single operation. This is itself a legitimate and actionable finding: on a dedicated Linux cloud instance, latency stayed within a narrow band across three independent runs; on a consumer Windows laptop under normal day-to-day load, the same operations varied by up to 4x run to run. An institution planning a Windows-hosted deployment should provision for the upper end of this range, not the lower end, and should expect materially more predictable latency on a dedicated Linux host. We report both runs rather than the faster or the more favorable one.

Across both environments and all measured runs, the cryptographic primitives - hashing and HMAC verification - stayed in the single-to-low-double-digit microsecond range, negligible relative to any underlying model call latency of tens to hundreds of milliseconds. The dominant and most variable cost is policy evaluation and receipt generation, both of which run at check-time against code changes or once per decision - not on a per-token hot path.

5. Regulatory Compliance Mapping

The following table maps Anchor's mechanisms to specific provisions of the five frameworks. Important limitation: this mapping has not been reviewed by external legal counsel or validated by any regulatory body. What Anchor provides is the technical substrate that makes a conformity assessment possible to perform; whether a given regulator's examiners would accept a given field as sufficient evidence of compliance with a given provision is a question this table cannot answer.

Framework	Provision	Requirement	Anchor Mechanism
EU AI Act	Art. 12	Automatic, tamper-evident logging	Decision Audit Chain; hash-chaining + HMAC
EU AI Act	Art. 13	Transparency / explainability	No-Prose Rule; structured reason codes
EU AI Act	Art. 9	Risk management system	Signed constitution.anchor; GOVERNANCE.lock
CFPB Reg. B	Adverse action	Reason specificity for AI credit decisions	adverse_action_reasons(); ReasonCode field
CFPB Reg. B / Circ. 2024-03	Explanation	No generic denial reasons	FeatureAttribution; structural enforcement
RBI FREE-AI	Rec. 7	Per-decision audit trail, CIMS-reportable	DAC; cims_payload()
RBI FREE-AI	Rec. 14	Explainable credit decisions	No-Prose Rule + FeatureAttribution
RBI FREE-AI	Rec. 19	Algorithmic fairness audits	ETH-001 prohibited-proxy taxonomy
SEC 2026	Audit trails	AI trading system audit trails	Full DAC; git_commit provenance

Framework	Provision	Requirement	Anchor Mechanism
SEC 2026	AI washing	Marketing claims match actual controls	GOVERNANCE.lock sealed rule set as evidence
NIST AI RMF	Govern 1.1	Documented governance policies	Signed constitution.anchor
NIST AI RMF	Manage 4.1	Continuous monitoring	AnchorRuntime continuous evaluation

6. Implementation Pathway

6.1 Integration Effort

For static (Layer 1) governance: pip install anchor-audit && anchor init --all && anchor check . - installs all domains and regulator files (~2 minutes) and runs the first scan immediately. For runtime (Layer 2) governance, import anchor.runtime auto-activates interception for calls made through registered SDK integrations (OpenAI, Anthropic, LangChain-based pipelines). High-stakes decision functions - credit, lending, trading - are wrapped with the @anchor.enforce decorator to require Structured Mode.

6.2 Deployment Models

For institutions with data residency requirements under EU GDPR, India's DPDP Act, or internal policies prohibiting transmission of source code or customer data externally, Anchor supports air-gapped initialization (anchor init --no-sign). The Decision Audit Chain writes exclusively to local disk regardless of configuration; no prompt, response, or audit data is transmitted over a network in the default audit-write path. The Sovereign Relay mechanism allows an authorized party to verify governance is operating correctly without receiving the underlying data.

6.3 Governance Assessment Structure

A governance assessment - running Anchor against an institution's existing AI-adjacent codebase, mapping findings to the relevant regulatory framework, and producing a findings brief - can be completed in three to five business days with read-only access and no changes to production systems.

7. Scope of Deterministic Guarantees

The claim is: for the class of violations represented in an organization's sealed constitution, and detectable by the corresponding mechanism in mitigation.anchor, Anchor's enforcement is deterministic - the rule fires if and only if the pattern is present, and (for BLOCKER-severity findings in CI) the commit is blocked if and only if the rule fires.

This guarantee has explicit boundaries. A developer can annotate a finding as an authorized exception - logged and individually attributed, but not blocking. A violation pattern not yet encoded in the mitigation catalog will not be detected by the corresponding static rule. A pattern existing only as a consequence of runtime control flow not visible in the AST falls to Layer 2's runtime interception.

> *The Correct Summary. Anchor deterministically enforces the violations its constitution defines. It does not claim to anticipate violations its constitution does not yet define.*

8. Comparison to Existing Approaches

Tool / Category	Core Question It Answers	Answers Authorization?
Microsoft Purview	What happened to this data, and who has access?	No - governs data, not AI decisions
Fiddler AI / Model Monitors	Is this model's output quality degrading?	No - governs model quality, not authority
Weights & Biases / MLOps	What happened during training and evaluation?	No - governs training, not execution
Guardrails AI / NeMo	Does this specific output conform to a schema?	Partially - schema only
Llama Guard / Content Classifiers	Is this content safe to surface to a user?	No - governs content, not authority
Open Policy Agent (OPA)	Does this request comply with defined policies?	Partially - policy enforcement, not AI-specific
Anchor	Was this action authorized, and can we prove it?	Yes - primary design goal

Anchor asks a question that comes after all of these have been satisfied: given that an output exists, and given that it may even be perfectly accurate, well-formed, and appropriate in content - is the system permitted to act on it, and can that permission be proven? This is why Anchor does not displace the tools above in a typical deployment; these are complementary layers of a stack, not competing implementations of the same layer.

9. Conclusion

The five automated financial system failures examined in this paper - costing more than $3 billion in aggregate - share a common structural pattern: the absence of infrastructure capable of enforcing, in real time and with a verifiable record, the boundary between what an automated system was authorized to do and what it actually did. Five regulatory frameworks have independently converged on the same response.

This requirement is satisfiable today. Not with a better classifier, not with more sophisticated safety training, but with well-understood software engineering and cryptographic techniques applied to the specific layer where the five failures above occurred: the layer between “the system produced an output” and “the system was authorized to act on it.”

The institutions that close the execution governance gap before an incident forces the question will be in a materially different position - both operationally and with their regulators - than those that close it afterward.

Aug 2026
EU AI Act Enforcement Begins

€30M / 6%
Maximum EU AI Act Penalty

$89M+
CFPB / Apple Card Enforcement (2024)

Anchor is available as open-source software under the Apache 2.0 license at github.com/AnimusLab/anchor and via pip install anchor-audit. Interactive policy evaluation is available at anchorgovernance.tech.

References

[1] European Parliament and Council. Regulation (EU) 2024/1689 (Artificial Intelligence Act). Official Journal of the European Union, 2024.

[2] Reserve Bank of India. Framework for Responsible and Ethical Enablement of AI (FREE-AI). Technical report, RBI, August 2025.

[3] Consumer Financial Protection Bureau. CFPB Orders Apple and Goldman Sachs to Pay Over $89 Million for Apple Card Failures. Enforcement action, October 23, 2024. (Dispute-handling / servicing failures under the CFPA and TILA - see Section 3.2 revision note.)

[4] Consumer Financial Protection Bureau. Circular 2024-03: Adverse action notification requirements and AI. CFPB Circular, 2024.

[5] Consumer Financial Protection Bureau. Statement of Director Rohit Chopra on algorithmic underwriting and adverse action notice requirements, 2023.

[6] New York State Department of Financial Services. Report on Apple Card Investigation. NYDFS, March 2021.

[7] U.S. Securities and Exchange Commission. 2026 Examination Priorities. SEC Division of Examinations, January 2026.

[8] National Institute of Standards and Technology. AI Risk Management Framework (AI RMF 1.0). NIST AI 100-1, January 2023.

[9] U.S. Securities and Exchange Commission. In the Matter of Knight Capital Americas LLC. Release No. 70694, File No. 3-15570, October 16, 2013.

[10] U.S. Securities and Exchange Commission and U.S. Commodity Futures Trading Commission. Findings Regarding the Market Events of May 6, 2010. Joint Report, September 30, 2010.

[11] U.S. Department of Justice. U.K. Trader Charged with Illegally Manipulating Stock Market. Press release, April 2015.

[12] Zillow Group, Inc. Form 10-Q for the Quarterly Period Ended September 30, 2021. SEC EDGAR, November 2, 2021.

[13] Financial Conduct Authority and Prudential Regulation Authority. TSB fined £48.65m for operational resilience failings. FCA Press Release, December 20, 2022.

[14] TSB Bank. Independent Review of TSB's IT Migration Programme. Slaughter and May, 2019.

[15] Citibank N.A. v. Brigade Capital Management, L.P. U.S. District Court, SDNY, No. 20-06539, February 16, 2021.

[16] Financial Conduct Authority. Artificial intelligence: How FCA principles apply. Discussion Paper DP 5/4, 2024.

[17] L. Kao. Constant-size cryptographic evidence structures for regulated AI. arXiv:2511.17118, 2025.

[18] B. Schneier and J. Kelsey. Secure audit logs to support computer forensics. ACM TISSEC, 2(2):159–176, 1999.

[19] S. Crosby and D. Wallach. Efficient data structures for tamper-evident logging. USENIX Security, 2009.

[20] M. Brunsfeld et al. Tree-sitter. tree-sitter.github.io, 2024.

[21] FINOS Foundation. Architecture as Code. github.com/finos/architecture-as-code, 2024.

[22] I. D. Raji et al. Closing the AI accountability gap. ACM FAccT, 2020.

About the Author

Tanishq Dasari is an independent researcher and founder of AnimusLab, an open-source AI governance infrastructure organization. AnimusLab is an active contributor to FINOS Labs and contributed to the DTCC Hackathon 2026, with a merged contribution implementing automated mapping from architectural components to governance mitigations within the FINOS reference architecture. Correspondence: tan@animuslab.dev. Anchor engine: github.com/AnimusLab/anchor. Research: animuslab.dev.