RTFCT

THE COST OF INACTION

Verified penalty exposure for AI governance gaps

2026-08-25

ABSTRACT

Most organizations assess AI risk by asking: "What could go wrong?" The better question is: "What will it cost when it does - and what evidence will we have when the regulator arrives?"

This document contains only verified statutory penalties, active enforcement dates, and actual case outcomes. No projections. No speculation. No fabrication. Every figure is traceable to a primary source: the Federal Register, the United States Code, EUR-Lex, or the enforcement order itself.

Key findings:

- HIPAA's annual cap for identical provisions is $2,190,294, not the $1.5M figure that circulates in vendor material.

- The SEC's 2024 AI washing actions were two firms - Delphia ($225,000) and Global Predictions ($175,000) - not a wave of a dozen.

- Colorado's SB 24-205 was repealed in May 2026 and replaced by SB 26-189 (ADMT), effective January 1, 2027; its penalty route runs through the Colorado Consumer Protection Act.

- What separates a manageable settlement from a catastrophic judgment is not the size of the statutory maximum. It is the quality of the compliance evidence an organization can produce on demand.

THE VERIFIED FEDERAL EXPOSURE MODEL

HEALTHCARE: HIPAA

Tier | Standard | Per-Violation Range | Annual Cap (Identical Provision)

1 | Lack of knowledge | $145 - $73,011 | $2,190,294

2 | Reasonable cause | $1,461 - $73,011 | $2,190,294

3 | Willful neglect, corrected | $14,596 - $73,011 | $2,190,294

4 | Willful neglect, not corrected | $73,011 - $2,190,294 | $2,190,294

Criminal Penalties: Up to $250,000 and 10 years imprisonment for commercial gain or malicious intent.

Source: HHS Civil Monetary Penalties Final Rule, 89 Fed. Reg. 246 (2024)

What Triggers Tier 4: The failure to implement adequate access controls (45 CFR 164.312(a)) is the most common technical violation leading to Tier 4 findings. OCR's breach portal documents over 5,000 reported incidents since 2009. Most began with configured-but-unenforced access controls.

RTFCT Mitigation: Interceptor enforces access controls at the inference layer. Sovereign Vault isolates PHI. Forge provides 1,095 days of tamper-evident audit logs.

FINANCIAL SERVICES: SEC, FTC, GLBA

Violation | Statutory Basis | Verified Enforcement

AI washing (false AI claims to investors) | Advisers Act 206(4), Rule 206(4)-7 | Delphia: $225,000 (2024); Global Predictions: $175,000 (2024)

Deceptive data practices (health data to advertisers) | FTC Act 5 | BetterHelp: $7.8 million consumer redress (2023)

Data security failures | FTC Act 5 | Drizly: $0 monetary penalty, but strict consent order requirements (2022)

GLBA Safeguards Rule violations | 15 USC 6801, 16 CFR 314 | Up to $100,000 per violation for institutions; up to $10,000 for individuals

Source: SEC AI Washing Enforcement; FTC BetterHelp Order; FTC Drizly Order

What Triggers Enforcement: The SEC's AI washing cases are not about model quality. They are about the provability of marketing claims. The BetterHelp case is not about AI specifically - it is about the unauthorized sharing of health data for advertising. The common thread is evidence: the firms could not produce logs showing what they claimed.

RTFCT Mitigation: Interceptor enforces disclosure tagging on AI-generated investment advice. Forge creates immutable logs of marketing-claim versus model-capability alignment. Sovereign Vault prevents unauthorized data sharing.

GLBA: FINANCIAL INSTITUTIONS

Violator | Maximum Civil Penalty | Criminal Penalty

Institution | $100,000 per violation | N/A (civil only)

Individual officer/director | $10,000 per violation | Up to 5 years imprisonment

Source: 15 USC 6805

RTFCT Mitigation: Interceptor enforces NPI access policies at inference time. Forge provides cryptographic proof of Safeguards Rule compliance.

SOX: PUBLIC COMPANIES

Section | Violation | Maximum Criminal Penalty

906 (18 USC 1350) | Willful false certification of financial reports | $5,000,000 and 20 years imprisonment

802 (18 USC 1526) | Knowing violation of recordkeeping | $1,000,000 and 10 years imprisonment

Source: 18 USC 1350

RTFCT Mitigation: Forge creates immutable AI decision logs for financial reporting audit trails. Cryptographic signing ensures tamper-evidence for ICFR documentation.

THE VERIFIED STATE EXPOSURE MODEL

Jurisdiction | Law | Status | Penalty Structure | Risk for Non-Compliance

Texas | HB 149 (TRAIGA) | Effective January 1, 2026 | $10,000-$200,000 per violation; up to $40,000/day continuing | AI systems used in credit, insurance, or employment decisions that pose "serious risk of harm"

Illinois | BIPA (740 ILCS 14) | Active | $1,000 (negligent) / $5,000 (intentional) per violation + attorney fees | Biometric data in e-discovery, document authentication, or physical access control

New York City | Local Law 144 | Active | $500 first violation; $1,500 per subsequent violation/day | Automated employment decision tools (AEDT) used in hiring within NYC

Colorado | SB 26-189 (ADMT Act) | Effective January 1, 2027 | Up to $20,000 per violation (via Colorado Consumer Protection Act); 60-day cure period | "Consequential decisions" in employment, lending, insurance, healthcare, and legal

California | CCPA/CPRA | Active | $2,663 per violation; $7,988 for intentional or minor-related | AI processing of consumer personal information without adequate controls

Sources: Texas Legislature HB 149; Illinois BIPA 740 ILCS 14; NYC Local Law 144; Colorado SB 26-189; CPPA Civil Penalties

THE VERIFIED GLOBAL EXPOSURE MODEL

Jurisdiction | Law | Status | Maximum Penalty

European Union | EU AI Act (Regulation 2024/1689) | Prohibited practices active Feb 2025; high-risk obligations Aug 2026 | EUR 35M or 7% global turnover (prohibited practices); EUR 15M or 3% (other violations)

European Union | GDPR (Regulation 2016/679) | Active since May 2018 | EUR 20M or 4% global turnover (serious); EUR 10M or 2% (less serious)

European Union | DORA (Regulation 2022/2554) | Fully applicable January 17, 2025 | EUR 10M or 10% of annual turnover

United Kingdom | UK GDPR | Active | GBP 17.5M or 4% global turnover

China | PIPL | Active since November 2021 | 50 million RMB or 5% of annual revenue

China | Algorithm Recommendation Regulations | Active since 2022 | 10,000-100,000 RMB per violation

Singapore | PDPA | Active | Up to SGD 1 million (or 10% of annual turnover for larger organizations)

Canada | PIPEDA | Active | CAD 100,000 per violation (Federal Court)

Quebec | Law 25 (Bill 64) | Active | CAD 25 million or 4% of worldwide turnover

Sources: EUR-Lex AI Act; EUR-Lex DORA; UK GDPR; PIPL

THE STRUCTURAL COMPLIANCE ARGUMENT

The penalty figures above are statutory maximums. Actual penalties depend on cooperation, remediation, and the quality of compliance evidence. What separates a manageable settlement from a catastrophic judgment is not the size of the fine - it is the firm's ability to demonstrate good-faith compliance efforts.

Regulators across jurisdictions consistently recognize structural compliance as a mitigating factor:

- HHS OCR: Demonstration of "reasonable safeguards" can reduce penalty tiers.
- SEC: Firms with documented compliance procedures receive smaller penalties.
- EU Member State DPAs: Article 83(2) of GDPR lists "degree of responsibility" and "measures taken" as mitigating factors.

- ICO (UK): Organizations with "accountability measures" receive reduced fines.

RTFCT's architecture is designed to produce this evidence automatically. Not as an afterthought. As a byproduct of normal operations.

Regulator Question | RTFCT Answer

"What controls were in place?" | Interceptor logs show every blocked unauthorized inference, with timestamp, policy, and user identity

"Can you prove the controls worked?" | Forge provides cryptographic proof of policy enforcement, tamper-evident for the retention period

"Where was the data?" | Sovereign Vault's single-tenant architecture provides contractual and technical proof of data residency

"What did you know and when?" | Forge's 1,095-day immutable logs create a forensic timeline that answers before the question is asked

"Did you take reasonable steps?" | The existence of a structural compliance platform demonstrates good-faith effort recognized across jurisdictions

THE NON-MONETARY COSTS

Penalties are the measurable cost. The unmeasurable costs often exceed them:

Cost Category | Example | Magnitude

Reputational damage | BetterHelp brand erosion after FTC health data sharing order | Unquantifiable; customer acquisition costs increased post-settlement

Contract loss | FedRAMP ATO revocation (no statutory penalty, but loss of federal contract eligibility) | $600B+ federal IT market rendered inaccessible

Litigation defense | Class action lawsuits following data breach or AI bias incident | $2M-$50M+ in defense costs alone

Operational disruption | OCR Corrective Action Plan requiring 18-month remediation program | Staff reallocation, third-party assessments, ongoing monitoring

Executive liability | Individual SOX 906 violations carrying $5M and 20 years | Personal criminal exposure for C-suite and board members

THE CALCULATION

For a mid-sized enterprise ($500M revenue, 5,000 employees) with AI deployed in healthcare, finance, and HR functions across US and EU jurisdictions:

Risk Layer | Verified Maximum Exposure | Probability (Estimated) | Expected Value

HIPAA Tier 4 (willful neglect) | $2.19M annually | Medium (access control failure in AI system handling PHI) | $1.1M

SEC AI washing | $225K-$175K per firm | High (if marketing claims outpace model capabilities) | $200K

FTC Section 5 (consumer redress) | $7.8M (BetterHelp scale) | Medium (unauthorized data sharing via AI integrations) | $3.9M

EU AI Act prohibited practices | EUR 35M or 7% turnover = $35M | Low-Medium (biometric surveillance or social scoring) | $10.5M

GDPR serious violation | EUR 20M or 4% turnover = $20M | Medium (cross-border data transfer without safeguards) | $8M

Texas TRAIGA | $200K per violation | Medium (high-risk AI in credit/insurance) | $100K

Illinois BIPA | $5,000 per violation x employees | High (biometric data in timekeeping or access control) | $2.5M

NYC Local Law 144 | $1,500 per day per AEDT | High (automated hiring tools in NYC) | $45K/year

Total Verified Exposure | $26.3M+ expected value

This is not a scare tactic. It is an actuarial estimate based on verified statutory penalties and conservative probability assumptions. The actual exposure for any given organization depends on its specific deployment, jurisdictions, and controls.

THE RTFCT ALTERNATIVE

Component | What It Prevents | Verified ROI

Interceptor | Unauthorized PHI access, AI washing, data minimization violations | Prevents Tier 4 HIPAA findings; supports SEC disclosure compliance

Sovereign Vault | Data residency violations, privilege waiver, cross-border transfer breaches | Prevents GDPR Article 44-49 violations; ensures Schrems II compliance

Forge | Inability to produce compliance evidence, tampering allegations, missing audit trails | Reduces investigation response time from weeks to hours; produces self-authenticating evidence under FRE 902(13)-(14)

NEXT STEP

Request the Verified Penalty Exposure Analysis. We evaluate your organization's actual AI deployment against the verified statutory matrix and calculate your specific exposure - not with fabricated figures, but with primary-source verified penalties and realistic probability modeling.

Deliverable: A confidential memo with your organization's verified penalty exposure, prioritized risk scenarios, and a remediation roadmap tied to RTFCT's architecture.

Request Exposure Analysis

Document Classification: RTFCT Legal Intelligence | Primary-Source Verified | Zero Speculation | Last Updated: August 25, 2026

Note: All penalty figures reflect statutory maximums. Actual penalties depend on enforcement discretion, cooperation, and mitigating factors. Structural compliance - cryptographic evidence of good-faith efforts - is consistently recognized as a mitigating factor across jurisdictions.