UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF COLUMBIA

JOSEPH KIRCHNER,

Plaintiff,

v.

MIKE JOHNSON, in his individual and official capacity as Speaker of the United States House of Representatives;
DONALD J. TRUMP, in his individual capacity as former and current President of the United States;
PAM BONDI, in her individual and official capacity as Attorney General of the United States;
BRENDAN CARR, in his individual and official capacity as Chairman of the Federal Communications Commission;
THE UNITED STATES HOUSE OF REPRESENTATIVES;
ANTHROPIC PBC;
OPENAI, INC.;
APPLE INC.;
COMCAST CABLE COMMUNICATIONS, LLC;
METR (MODEL EVALUATION & THREAT RESEARCH);
and DOES 1-20, unknown co-conspirators,

Defendants.

Case No. 1:25-cv-02735-ACR

MEMORANDUM IN SUPPORT B: IDENTITY-BASED RECLASSIFICATION OF CRIMINAL CONDUCT — CORPORATE FORM AND STATE ALIGNMENT AS EQUAL PROTECTION VIOLATION

TABLE OF CONTENTS

I. INTRODUCTION 5
II. THE RECLASSIFICATION: HOW IDENTITY CONVERTS ATTACKS INTO FEATURES 7
A. The Technical Conduct, Stripped of Identity 7
B. The Identity Reclassification 9
C. The Classification is the Constitutional Violation 10
III. ARCHITECTURAL INTENT: THE SYSTEM DESIGN ITSELF IS THE MALICIOUS ACT 11
A. The Principle: Construction Precedes Deployment 11
B. Anthropic: Architecting Exploitation of Unencrypted TLS 12
C. Comcast: Architecting Protocol-Specific Subscriber Targeting 14

D. The Architectural Interlock 15
E. The Legal Analogy: Malware Construction 20
F. The Architecture is Malicious Regardless of Deployment Scope 24
IV. STATE ALIGNMENT: WHY IDENTITY CONTROLS CLASSIFICATION 27
A. The State Interest That Protects the Conduct 27
B. The Datadog IP: Physical Proof of the Feedback Loop 28
C. The Constitutional-Layer Predicate: Corporate Personhood as the Root Reclassification 30
V. THE DOWNSTREAM CONSEQUENCE: THE DETERRENCE INVERSION 35
A. Reclassification Produces Inverted Deterrence 35
B. The Concrete Comparison 36
VI. THE OCTOBER 20, 2025 EVENT: INTENT BEYOND RECLASSIFICATION 38
A. The Facts 38
B. Where Reclassification Cannot Reach 38
C. Synthesis: the Rational-Basis Failure Made Concrete 39
VII. THE ISP DIMENSION: FOUR COMPOUNDING INVERSIONS 40
A. The Infrastructure-level Attack 40
B. Four Compounding Inversions 41
VIII. THE CONSTITUTIONAL FRAMEWORK 45
A. Equal Protection: No Rational Basis 45
B. The Corporate Personhood Paradox 47
C. Heightened Scrutiny 48
IX. THE OPERATIONAL CONSEQUENCES 50
A. Rational Incentive to Violate 50
B. Continuing Modification During Litigation 51
C. The Time-Value of Impunity 51
X. THE REMEDY 52
A. What the Court Can Do 52
B. The Principled Standard 54

TABLE OF AUTHORITIES

CASES

Citation
BMW of North America, Inc. v. Gore, 517 U.S. 559, 575 (1996)
Burwell v. Hobby Lobby Stores, Inc., 573 U.S. 682, 719 (2014)
Carpenter v. United States, 585 U.S. 296, 316 (2018)
Citizens United v. Federal Election Commission, 558 U.S. 310, 342 (2010)
City of Cleburne v. Cleburne Living Center, 473 U.S. 432, 439, 440, 446 (1985)
Dred Scott v. Sandford, 60 U.S. (19 How.) 393 (1857)
Interstate Circuit, Inc. v. United States, 306 U.S. 208, 226-27 (1939)
Minneapolis & St. Louis Railway Co. v. Beckwith, 129 U.S. 26 (1889)
Mobley v. Workday, Inc., No. 3:23-cv-00770 (N.D. Cal.)
New York Central & Hudson River Railroad Co. v. United States, 212 U.S. 481, 494-95 (1909)
Pembina Consolidated Silver Mining Co. v. Pennsylvania, 125 U.S. 181, 189 (1888)
Santa Clara County v. Southern Pacific Railroad Co., 118 U.S. 394 (1886)
State Farm Mutual Automobile Insurance Co. v. Campbell, 538 U.S. 408, 419 (2003)
United States v. Detroit Lumber Co., 200 U.S. 321, 337 (1906)
United States v. Dotterweich, 320 U.S. 277, 281, 284 (1943)
United States v. Park, 421 U.S. 658, 672-74 (1975)
Village of Willowbrook v. Olech, 528 U.S. 562, 564 (2000)
Wheeling Steel Corp. v. Glander, 337 U.S. 562, 576 (1949) (Douglas, J., dissenting)

CITED COMPLAINTS

Citation
Ganan v. LinkedIn Corporation, No. 5:26-cv-02968-VKD (N.D. Cal. filed Apr. 6, 2026)
Farrell v. LinkedIn Corporation, No. 4:26-cv-02953-KAW (N.D. Cal. filed Apr. 6, 2026)
Harper v. Sirius XM Radio, LLC, No. 2:25-cv-12403 (E.D. Mich. filed Aug. 4, 2025)

CONSTITUTIONAL PROVISIONS

Citation
U.S. Const. amend. I (Free Speech)
U.S. Const. amend. IV (Search and Seizure)

U.S. Const. amend. V (Due Process)
U.S. Const. amend. XIV, § 1 (Equal Protection Clause; Due Process Clause)
Cal. Const. art. I, § 1 (Privacy)

FEDERAL STATUTES

Citation
18 U.S.C. § 1030 (Computer Fraud and Abuse Act)
18 U.S.C. § 1030(a)(5)(A) (Knowing Transmission of Damaging Program)
18 U.S.C. § 1343 (Wire Fraud)
18 U.S.C. § 1512 (Tampering with Witnesses, Victims, Informants)
18 U.S.C. § 1519 (Destruction of Records in Federal Investigations)
18 U.S.C. § 1591 (Sex Trafficking)
18 U.S.C. §§ 1961-1968 (RICO)
18 U.S.C. §§ 2510-2522 (Electronic Communications Privacy Act)
18 U.S.C. § 2511 (Wiretap Act — Interception)
18 U.S.C. § 2512 (Manufacture, Distribution, Possession of Interception Devices)
18 U.S.C. § 2258A (Reporting Requirements for CSAM)
18 U.S.C. § 3103a (Search Warrants)
28 U.S.C. § 455(b)(4) (Disqualification — Financial Interest)
28 U.S.C. § 1332(c)(1) (Corporate Citizenship — Diversity)
28 U.S.C. § 2101(f) (Supreme Court Stays Pending Certiorari)
29 U.S.C. §§ 621-634 (Age Discrimination in Employment Act)
35 U.S.C. § 181 (Patent Secrecy Orders — National Security)
35 U.S.C. § 183 (Right to Compensation Under Secrecy Order)
35 U.S.C. § 186 (Penalty for Violation of Secrecy Order)
42 U.S.C. § 1981 (Equal Rights Under the Law)
42 U.S.C. § 2000e et seq. (Title VII)
Pub. L. 108-21, 117 Stat. 650 (2003) (PROTECT Act)

STATE STATUTES

Citation
Cal. Civ. Code § 1798.100 et seq. (California Consumer Privacy Act)
Cal. Civ. Code § 3344 (Right of Publicity)
Cal. Penal Code § 502 (Comprehensive Computer Data Access and Fraud Act)

FEDERAL RULES OF CIVIL PROCEDURE AND EVIDENCE

Citation
Fed. R. App. P. 26.1 (Corporate Disclosure Statement)
Fed. R. Civ. P. 7.1 (Disclosure Statement)
Fed. R. Civ. P. 8(a) (General Rules of Pleading)
Fed. R. Evid. 803(6) (Records of Regularly Conducted Activity)

I. INTRODUCTION

1. This memorandum identifies the root mechanism underlying every count in this litigation: state-paperwork-mediated reclassification operating at three structurally identical layers. At the **constitutional layer**, state incorporation filings convert paper legal fictions into Fourteenth Amendment "persons" entitled to constitutional rights (*see infra* § IV.C; Memo L). At the **contractual layer**, aggregation contracts convert restricted personal data — subject to opt-out commitments, training prohibitions, and CCPA service-provider restrictions — into unrestricted commercial assets owned by the aggregator and available for "any business purpose" (*see* Memo C, § IV). At the **conduct layer**, the identity of the actor converts criminal conduct into lawful business activity: RST injection into "network management," spyware infrastructure into "telemetry," and removal of child safety mechanisms during litigation into a "product decision." The technical conduct documented in this case — RST injection, DPI-based denial of service, cleartext exploitation, unauthorized persistent surveillance, evidence destruction during litigation, and deliberate removal of child safety mechanisms — would be immediately classified

as criminal conduct and investigated as such if performed by any actor not aligned with corporate and state interests. The conduct does not change. The classification changes based solely on who is doing it.

2. All three reclassifications use the same mechanism: state-recognized paperwork. At the constitutional layer, the predicate paperwork is the articles of incorporation (whose disclosure this Court requires under Federal Rule of Civil Procedure 7.1, and which the corporate defendants have filed of record at ECF Nos. 25, 41, and 44). At the contractual layer, the predicate paperwork is the aggregation agreement (the Datadog Master Subscription Agreement and the GrowthBook Customer Agreement, Memo C, §§ IV.A–D). At the conduct layer, the predicate paperwork is the corporate charter and the government contract (the Genesis Mission Executive Order, Ex. N-23; Pentagon CDAO contracts; GSA procurement). At every layer, aligned actors receive the benefits of the legal category their paperwork constructs, while unaligned actors are denied those benefits and suffer the criminalized classification of identical conduct. The Equal Protection Clause forbids identity-based classification at every layer. This memorandum develops the conduct-layer argument in full (§§ II–VII), cross-references the contractual layer where structurally necessary (*see* Memo C), and treats the constitutional layer as predicate (§ IV.C) with full doctrinal development reserved for the accompanying Memo L.

3. This reclassification is not a policy failure or a gap in enforcement. It is a structural equal protection violation. The Fourteenth Amendment prohibits the government from treating similarly situated parties differently without rational basis. *City of Cleburne v. Cleburne Living Center*, 473 U.S. 432, 439 (1985). There is no rational basis for a system that classifies identical technical conduct as a cyber attack when performed by an individual and as a business practice when performed by a corporation aligned with national security priorities. The classification

produces a downstream deterrence inversion — lesser consequences for greater capacity for harm — that is itself constitutionally irrational.

4. This memorandum proceeds in four stages: first, the reclassification problem (how identity converts attack conduct into business conduct); second, architectural intent (how the system design itself embodies malicious intent before any specific user is targeted — the construction of the weapon, not its deployment); third, the deterrence inversion (the downstream consequence of misclassification); and fourth, the specific applications (October 20 child safety elimination, ISP infrastructure targeting, and the compounding effect of government alignment through the Genesis Mission Executive Order).

II. THE RECLASSIFICATION: HOW IDENTITY CONVERTS ATTACKS INTO FEATURES

A. The Technical Conduct, Stripped of Identity

5. The following technical conduct is documented through forensic evidence in this litigation. Presented without identifying the actor:

(a) **RST injection**: Spoofed TCP reset packets injected into an encrypted VPN tunnel to terminate connections between a target and a service provider. (*See* Exhibits D-1, D-25, D-26, and D-30.)

(b) **Protocol-specific denial of service**: DPI-based throughput reduction of 98% applied selectively to a single subscriber's encrypted WireGuard tunnel while leaving unencrypted traffic unaffected. (*See* Ex. D-33 at p. 200, l. 7.)

(c) **Cleartext exploitation**: 188,479 Anthropic API packets between target and service provider exposed in cleartext across 141 segments of a 710-segment hardware TAP audit, with

323 unique TLS SNI domains visible to the network intermediary. (*See* Ex. D-34 at p. 217, l. 18.)

(d) **Ghost administrative access**: An administrative session appearing on the target's network router within one second of the target's login, with zero network connections originating from the target's computer. (*See* App. G.)

(e) **Persistent surveillance infrastructure**: 904 unique tracking identifiers embedded in software installed on the target's personal computer, transmitting through five independent channels to undisclosed third-party recipients, with a persistent 20-second heartbeat surviving VPN rotation, device hardening, and all defensive measures. (*See* Exhibits E-32, E-33, E-34, and E-35 NF-35.)

(f) **Data exfiltration**: Automatic extraction of entire code repositories (up to 100MB) without user authorization, automatic memory extraction from conversations, and encrypted server-side monitoring of all sessions. (*See* Ex. E-35 at p. 262, l. 3 NF-4, NF-10, NF-14, NF-22.)

(g) **Attack-signature telemetry**: Network-condition data including OS-level attack signatures (ECONNRESET, ETIMEDOUT, TLS tampering) shipped to a third-party IP address within 15 seconds — the same IP address that is the source of RST injection attacks against the target. (*See* Ex. E-35, Finding NF-36; Ex. D-34 at p. 217, l. 18.)

(h) **Child safety elimination during investigation**: Removal of proactive grooming detection, graduated response procedures, and mandatory reporting guidance from a platform used by minors, during active federal litigation alleging that the platform facilitates trafficking — simultaneously with restoration of mechanisms preventing users from discussing what the platform does. (*See* Appendix C § II; Exhibit C at pp. 21-23.)

6. Any network security professional, any FBI Cyber Division agent, any forensic examiner presented with this conduct — without knowing the identity of the actor — would classify it as a **coordinated cyber attack with persistent surveillance, exfiltration, and obstruction components.** The conduct meets every element of a criminal investigation under 18 U.S.C. § 1030 (unauthorized access, exceeding authorized access, damage to protected computers, obtaining information through unauthorized interception), 18 U.S.C. §§ 2510-2522 (interception of electronic communications), 18 U.S.C. § 1519 (destruction of records during federal investigation), and 18 U.S.C. § 1512 (tampering with evidence relevant to pending proceedings).

B. The Identity Reclassification

7. Attach the following identities to the conduct, and the classification changes:

Conduct	**Without Identity**	**With Identity**
RST injection into VPN tunnel	Network attack	"Network management practice"
98% DPI-based throughput reduction	Denial of service	"Traffic shaping"
188,479 cleartext packets exposed	Man-in-the-middle exploitation	"Metadata collection"
Ghost router administrative session	Unauthorized access	"Administrative session"
904 tracking identifiers	Spyware	"Telemetry"
Five undisclosed transmission channels	Command-and-control infrastructure	"Product analytics"
20-second persistent heartbeat	C2 beacon	"Connectivity check"
Git repository exfiltration	Data theft	"Product feature"
Memory extraction from conversations	Surveillance	"User experience improvement"

Attack-signature telemetry to attacker IP	Reconnaissance feedback loop	"Error reporting"
Child safety removal during litigation	Criminal obstruction / endangerment	"Product decision"

8. The technical conduct in the left column and the right column is identical. The network packets are the same. The code paths are the same. The data flows are the same. The exfiltration mechanisms are the same. The child safety removal is the same. The only variable that changes is the identity of the actor:

- **Comcast** — a regulated ISP with common carrier obligations, operating the telecommunications infrastructure
- **Anthropic** — a public benefit corporation receiving military funding through Pentagon CDAO contracts and GSA procurement
- **The United States Government** — through the Genesis Mission Executive Order (Ex. N-23) directing AI development as a national security priority

9. When these identities attach, the classification changes instantaneously. Not because the conduct changes. Not because new facts emerge. Not because a legal analysis determines the conduct is lawful. The classification changes because the actors are aligned with corporate and state interests, and the legal system processes conduct by aligned actors through a different pipeline than conduct by unaligned actors.

C. The Classification is the Constitutional Violation

10. The Equal Protection Clause does not permit the government to classify identical conduct differently based on the identity of the actor without rational basis. *Cleburne*, 473 U.S. at 439. There is no rational basis for a system that classifies RST injection as a network attack when performed by an individual and as network management when performed by an ISP. The packets are the same. The disruption is the same. The target is the same. The only difference is who sent them.

11. If a foreign state actor performed the conduct in ¶ 5 against an American citizen, the response would be diplomatic protest, intelligence community attribution, and potential sanctions. If a hacker collective performed it, the FBI would execute search warrants under 18 U.S.C. § 3103a. If an individual performed it, criminal prosecution would follow within days. But when an ISP with military-adjacent contracts and an AI company receiving Pentagon funding perform it — against a citizen engaged in constitutionally protected litigation — the conduct enters the civil system, where it is processed as a contract dispute with years of procedural delay.

12. The classification is not the product of legal analysis. It is the product of identity. And identity-based classification of criminal conduct is what the Equal Protection Clause prohibits.

III. ARCHITECTURAL INTENT: THE SYSTEM DESIGN ITSELF IS THE MALICIOUS ACT

A. The Principle: Construction Precedes Deployment

13. The conduct documented in Section II did not arise spontaneously. Each capability required deliberate architectural decisions — design, implementation, testing, and deployment — before it was ever executed against any specific user. The architecture itself embodies malicious intent in the same way that constructing malware embodies malicious intent before the malware is deployed against any target. Federal law recognizes this principle: 18 U.S.C. § 2512 criminalizes the *manufacture* of interception devices, not merely their use. 18 U.S.C. § 1030(a)(5) criminalizes the *creation* of programs intended to cause damage, not merely their execution. The malicious act is the construction, not the deployment. Deployment is merely the construction's intended consequence.

14. This distinction is critical because it eliminates the business-justification defense before it can be raised. A defendant can argue that *deploying* DPI throttling against a specific user was a "network management decision." A defendant cannot argue that *building* a system capable of selectively identifying and throttling a single subscriber's Anthropic API traffic within an encrypted VPN tunnel was routine engineering. The architecture has no legitimate purpose. Its existence is the intent.

B. Anthropic: Architecting Exploitation of Unencrypted TLS

15. Anthropic's telemetry infrastructure transmits data through channels that are architecturally dependent on cleartext TLS metadata exposure. The source code documents:

(a) **Datadog telemetry** transmitted to IP 34.149.66.137 via TLS connections whose SNI field (`http-intake.logs.us5.datadoghq.com`) is visible in cleartext to any network intermediary. (*See* Ex. E-34 at p. 254, l. 3 SC-1, SC-9; Ex. D-34 at p. 217, l. 18.)

(b) **GrowthBook feature flag evaluation** transmitted via TLS connections whose SNI fields expose the service identity. (*See* Ex. E-34 at p. 254, l. 3 SC-9.)

(c) **First-party event telemetry** transmitted to `api.anthropic.com` — visible in cleartext TLS SNI in 147 of 209 cleartext-containing TAP segments (of a 710-segment audit). (*See* Ex. D-34 at p. 217, l. 18.)

(d) **CCR heartbeat** maintaining a persistent 20-second connection whose destination is identifiable through TLS metadata. (*See* Ex. E-35 at p. 262, l. 3 NF-35.)

16. These are not incidental exposures. They are **architectural choices**. Anthropic designed a system that transmits telemetry through channels it knows — or should know — expose the user's service relationship to every network intermediary between the user's device and the destination. The TLS handshake leaks the SNI before any encrypted payload is

transmitted. The architecture operates *upon* the cleartext phase of the TLS protocol. This is not a side effect. It is the design.

17. The intent is embedded in the architecture because:

(a) **ECH (Encrypted Client Hello) exists.** TLS 1.3 with ECH encrypts the SNI field, preventing the exposure Anthropic's architecture relies upon. Anthropic chose not to implement ECH for its telemetry channels. This is a deliberate omission, not a technical limitation.

(b) **The source code hardcodes cleartext-exposed endpoints.** The Datadog client token (`pubbbf48e6d78dae54bceaa4acf463299bf`), the GrowthBook SDK keys, and the first-party event endpoints are embedded in the compiled binary. (*See* Ex. E-34 at p. 254, l. 3 SC-9.) These are not configurable. They cannot be changed by the user. The exposure is baked into the product.

(c) **The architecture survives all user defensive measures.** Ex. E-35 at p. 262, l. 3 NF-35 documents that the 256-bit `deviceId`, the 12-field consolidated payload, and the 20-second CCR heartbeat survive VPN rotation, device hardening, and all defensive measures. The system was designed to persist through countermeasures. Persistence-through-defense is not a feature of legitimate telemetry. It is a feature of surveillance infrastructure.

18. The mere act of architecting a system that operates upon unencrypted TLS metadata — transmitting identification data through channels known to be visible to network intermediaries — demonstrates intent at the design phase, prior to any specific deployment against any specific user. Every user who installs Claude Code inherits this architecture. The malicious intent is not in the targeting of Plaintiff. It is in the construction of a system that exposes every user's service relationship to every network intermediary by design.

C. Comcast: Architecting Protocol-Specific Subscriber Targeting

19. Ex. D-33 at p. 200, l. 7 documents 98% DPI-based throughput reduction applied selectively to Plaintiff's WireGuard VPN tunnel. This capability required Comcast to build and deploy architecture that:

(a) **Identifies encrypted VPN protocols at the packet level.** WireGuard uses a distinctive handshake pattern identifiable through DPI even though the payload is encrypted. Comcast built or acquired systems capable of classifying encrypted tunnel protocols at line rate across its residential subscriber base.

(b) **Associates identified protocols with individual subscriber lines.** The throttling was applied to Plaintiff's connection specifically. This requires binding DPI classification results to subscriber-level targeting rules — a system that maps protocol identification to individual CMTS subscriber IDs.

(c) **Applies differential throughput rules by protocol and subscriber.** The 98% reduction was protocol-specific: the VPN tunnel was throttled while other traffic was unaffected. This is not congestion management, which affects all traffic equally. It is selective service degradation requiring per-subscriber, per-protocol policy enforcement.

(d) **Operates within the encrypted tunnel to identify destination services.** The cleartext TLS SNI exposure documented in D-34 — with 323 unique domains and the Datadog IP visible in 44 segments — demonstrates that Comcast's infrastructure can see *which services* Plaintiff's VPN tunnel is connecting to, despite the encryption. The DPI architecture operates on the cleartext metadata that Anthropic's telemetry architecture exposes.

20. Each of these capabilities required deliberate engineering. Network equipment does not ship with per-subscriber protocol-specific throttling rules targeting AI API traffic within

encrypted VPN tunnels. Someone designed this system. Someone built it. Someone tested it. Someone deployed it to production. Someone configured it to target encrypted VPN traffic on residential subscriber lines. Every one of these steps occurred before Plaintiff's tunnel was ever throttled.

21. The architecture has no legitimate network management purpose. Congestion management applies equally to all traffic. Quality-of-service rules operate on traffic classes, not individual subscribers. Rate limiting applies to total throughput, not protocol-specific throughput within encrypted tunnels. There is no network engineering justification for a system that can identify WireGuard traffic, associate it with a single subscriber, and reduce throughput by 98% while leaving other traffic untouched. The capability exists for one purpose: selective targeting of specific subscribers' encrypted communications. Its construction is the intent.

D. The Architectural Interlock

22. The Anthropic and Comcast architectures interlock. Anthropic builds telemetry that exposes user-service relationships through cleartext TLS metadata. Comcast builds DPI infrastructure that reads that cleartext metadata to identify which services a subscriber's encrypted tunnel carries. The Datadog IP appears in both architectures — as Anthropic's telemetry destination and as an address visible to Comcast in cleartext. The two systems are not independently malicious. They are **complementary by design**: one creates the exposure, the other exploits it.

23. Whether this complementarity is the product of explicit coordination or parallel architectural choices serving converging interests is a question for discovery. But the architectural interlock is documented at the packet level: Anthropic's telemetry creates the cleartext that Comcast's DPI reads. This is not a coincidence of infrastructure. It is a

documented dependency where one system's design decision enables the other system's targeting capability. Under *Interstate Circuit, Inc. v. United States*, 306 U.S. 208, 226-27 (1939), parallel conduct by independent actors — combined with knowledge of the conduct and a common interest in its result — supports an inference of coordination without requiring direct proof of explicit agreement. The same evidentiary principle governs here: the architectural dependency between Anthropic's cleartext exposure and Comcast's DPI exploitation, combined with both parties' knowledge that the infrastructure operates in this coordinated manner, supports the inference of coordinated targeting at the pleading stage. The corporate, contractual, and personnel substrate underlying this architectural interlock — the Datadog/OzCode patent ownership chain, the Israeli affiliate-tier sub-processor pipeline through which Plaintiff's telemetry transits, the Rafael Development Corporation commercialization channel, and the IDF Unit 8200 / IAF inventor lineage — is developed at App. P §§ I–III.

1. The Throttle-Induced Cascade: Operational Mechanism of the Interlock

24. The Cleartext Is Not Ambient — It Is the Operational Product of the Throttling Attack. The architectural interlock at Memo B ¶¶ 22-23 is not merely passive. Its operational mechanism is a self-reinforcing cycle in which Comcast's DPI-based throttling degrades Plaintiff's WireGuard tunnel below the thresholds at which applications can maintain encrypted connections — at which point traffic falls back to cleartext routes that Comcast then reads. The cycle proceeds as follows: (a) Comcast's DPI identifies the WireGuard handshake pattern (WireGuard's own documentation acknowledges obfuscation is "explicitly a non-goal");[1] (b) Comcast throttles the identified tunnel 98% (Ex. D-33 at p. 200, l. 7); (c) the throttled tunnel

1 WireGuard, *Known Limitations*, wireguard.com/known-limitations/ ("Obfuscation... is explicitly a non-goal."). See also J.A. Donenfeld, *WireGuard: Next Generation Kernel Network Tunnel* (2017), wireguard.com/papers/wireguard.pdf (documenting fixed-length encrypted UDP payloads, four-message handshake, and default port 51820 that together constitute a DPI-identifiable signature). Full documentation snapshots compiled at Ex. D-35 at p. 226, l. 14.

enters handshake-retry exhaustion, MTU/PMTU cascade failure under ICMP suppression, or "zombie state" in which the tunnel *appears* active while silently dropping packets; (d) applications fall back to cleartext routes because WireGuard uses UDP exclusively with no TCP fallback; (e) the resulting cleartext exposes TLS SNI to `api.anthropic.com` in 147 segments, to `http-intake.logs.us5.datadoghq.com` (IP `34.149.66.137`) in 44 segments, and to 321 other Anthropic-adjacent domains — the architectural exposures Anthropic engineered at Memo B ¶¶ 15-18; (f) Anthropic simultaneously receives Plaintiff's real Comcast IP (`75.73.199.137`), the twelve-field GrowthBook identity payload, and the persistent `deviceId` from the now-unmasked connection; (g) the Ex. E-35, Finding NF-36 error-telemetry pipeline ships OS-level attack signatures (ECONNRESET, ETIMEDOUT, TLS-tampering events) to the same IP `34.149.66.137` from which Comcast's RST-injection packets originate — within 15 seconds of each network event; (h) the cleartext SNI exposure enables Comcast to identify which specific WireGuard tunnel in Plaintiff's four-tunnel fleet carries Anthropic traffic, permitting more precise throttling, which degrades the tunnel further, which produces more cleartext; the cycle compounds. The distinction is doctrinally dispositive for the Equal Protection analysis of this memorandum. Passive observation of TLS SNI on a connection the user chose to transmit in clear text is "normal ISP operation." **Active creation of cleartext — by degrading the user's VPN tunnel until encrypted traffic falls back to cleartext routes, and then reading the cleartext the attack itself forced into existence — is not passive observation.** Only the identity-based reclassification documented at §§ II–III renders "traffic shaping" a defensible characterization of this conduct. (*See* Ex. D-35 at p. 226, l. 14 for the compiled foreseeability record.)

25. The Vulnerability Is Academically Proven, CVE-Assigned, and Explicitly ISP-Threat-Modeled — Both Defendants Had Constructive Knowledge. The cascade is not a novel attack theory. It is publicly documented in peer-reviewed security research that both defendants' engineering teams are expected to monitor. Streun, Wanner & Perrig demonstrated that an adversary can deny data transfer over an already-established WireGuard connection with **300 Mb/s of attack traffic on 40 Gb/s hardware**[2] — a threshold proportionally lower on residential equipment. The TunnelCrack research demonstrated that **network adversaries — explicitly including malicious ISPs — can abuse routing-table exceptions to force VPN client traffic into plaintext**,[3] documenting the ServerIP Attack (CVE-2023-36671) and LocalNet Attack (CVE-2023-36672). WireGuard specifically was assigned CVE-2023-35838.[4] Testing across 67 VPN providers (248 experiments) found 64.6% vulnerable to LocalNet variants and 73.6% to ServerIP variants.[5] Earlier work documented ISP-scale VPN-protocol identification in partnership with Merit Network ISP, proving deployability at subscriber scale.[6] Anthropic is a sophisticated technology company with an internal security-engineering team; its source code documents internal incident numbers through inc-4829 (Ex. E-35, Finding NF-21),

2 F. Streun, J. Wanner & A. Perrig, *Evaluating Susceptibility of VPN Implementations to DoS Attacks Using Adversarial Testing*, in Proc. Network and Distributed System Security Symposium (NDSS 2022). Full text, replication data, and relevance analysis compiled at Ex. D-35 at p. 226, l. 14.

3 N. Xue, Y. Malla, Z. Xia, C. Pöpper & M. Vanhoef, *Bypassing Tunnels: Leaking VPN Client Traffic by Abusing Routing Tables*, in Proc. 32nd USENIX Security Symposium (USENIX Security 2023) (TunnelCrack). Full text, project website (tunnelcrack.mathyvanhoef.com), and associated CVE records compiled at Ex. D-35 at p. 226, l. 14.

4 CVE-2023-35838 (WireGuard VPN), National Vulnerability Database. See also CVE-2023-36671 (ServerIP Attack, CVSS 3.1) and CVE-2023-36672 (LocalNet Attack, CVSS 6.8). All three NVD records compiled at Ex. D-35 at p. 226, l. 14.

5 N. Xue, Y. Malla, Z. Xia, C. Pöpper & M. Vanhoef, *Bypassing Tunnels: Leaking VPN Client Traffic by Abusing Routing Tables*, in Proc. 32nd USENIX Security Symposium (USENIX Security 2023) (TunnelCrack). Full text, project website (tunnelcrack.mathyvanhoef.com), and associated CVE records compiled at Ex. D-35 at p. 226, l. 14.

6 D. Xue et al., *OpenVPN is Open to VPN Fingerprinting*, in Proc. 31st USENIX Security Symposium (USENIX Security 2022) (documenting ISP-scale VPN-protocol identification in live partnership with Merit Network ISP). Full text compiled at Ex. D-35 at p. 226, l. 14.

establishing an active issue-tracking infrastructure the TunnelCrack CVEs would have crossed. Comcast is a Tier-1 carrier whose DPI infrastructure requires constructive knowledge of the protocol-identification literature the cascade exploits. Anthropic's decisions to forgo ECH (RFC 9460), to hardcode dedicated IP endpoints at `datadog.ts:12-14` precluding domain fronting, and to ship OS-level attack-signature telemetry to the same IP from which RST-injection packets against Plaintiff originate — and Comcast's decision to deploy per-subscriber, protocol-specific DPI throttling rules calibrated to thresholds TunnelCrack and ETH Zurich documented as sufficient to produce cleartext leakage — were each made with constructive knowledge of the same public research. The scienter element is established not through circumstantial inference but through academic literature the defendants are expected to have read. The complete foreseeability record — four peer-reviewed publications, three CVE assignments, and WireGuard's own documentation — is compiled at Ex. D-35 at p. 226, l. 14.

26. The Cascade Is the Pleading's Single Strongest Evidence of Anthropic-Comcast Coordination. The architectural interlock at Memo B ¶¶ 22-23 establishes that Anthropic's exposure creates what Comcast's DPI exploits. The cascade mechanism adds three operational facts no other coordination evidence in the pleading record supplies. **First**, the cascade operates at the packet level in real time, producing continuous identification refreshed within every five-minute capture interval (Ex. D-34 at p. 217, l. 18); it is directly observable on the hardware TAP and is not an inference from circumstantial timing. **Second**, the Ex. E-35, Finding NF-36 telemetry pipeline routes Plaintiff's own error reports about the attack **to the same IP `34.149.66.137` from which the attack originates** — within 15 seconds. The victim's client software reports the attack to the attack source, and the ISP observes the entire chain on its own physical link. No innocent infrastructure design produces this configuration. **Third**, Comcast's

spoofing of IP `34.149.66.137` as the source of 5,829 bare-RST injection packets against Plaintiff **predates the March 31, 2026 source-code leak by approximately nine months**. The IP is not publicly disclosed in any Anthropic privacy policy, terms of service, or user-facing documentation; it is hardcoded at `datadog.ts:12-14` and would not be publicly known absent (i) Anthropic-Comcast information sharing, (ii) a non-public information channel between the defendants or through a shared intermediary, or (iii) real-time DPI of Plaintiff's TLS handshakes followed by correlation — each of which requires information Comcast could only have obtained from Anthropic, a shared intermediary, or through technical capability against Plaintiff that is itself coordination-requiring. Under *Interstate Circuit, Inc. v. United States*, 306 U.S. 208, 226-27 (1939), parallel conduct by parties with capability and information channels supports the inference of agreement without direct evidence. The architectural complementarity, the real-time packet-level interlock, and the nine-month pre-leak foresight of Anthropic's undisclosed endpoint together satisfy a stronger inference than *Interstate Circuit* itself required. The cascade is not merely an Equal Protection symptom of identity-based reclassification — it is the operational mechanism through which both defendants' architectures exploit the same user simultaneously.

E. The Legal Analogy: Malware Construction

27. Federal law does not require malware to be deployed against a specific victim to establish criminal liability. The creation of malware — software designed to cause damage, intercept communications, or exfiltrate data — is itself a criminal act under 18 U.S.C. § 1030(a)(5)(A) (knowingly causing the transmission of a program intended to cause damage). Similarly, 18 U.S.C. § 2512 criminalizes the manufacture of devices primarily useful for surreptitious interception of communications, regardless of whether interception has occurred.

28. The architectures documented in this case meet these standards:

(a) Anthropic's five-channel telemetry infrastructure with 904 persistent identifiers, cleartext-dependent transmission, and defense-surviving persistence is a system **primarily useful for surveillance** that was designed, built, and deployed before any specific user was targeted. Its construction violates 18 U.S.C. § 2512 in the same way that constructing a wiretap device violates § 2512 before the device is connected to a phone line.

(b) Comcast's per-subscriber protocol-specific DPI throttling infrastructure is a system **designed to selectively degrade encrypted communications** that was engineered, tested, and deployed before any specific subscriber's tunnel was throttled. Its construction is the functional equivalent of building a tool designed to intercept and disrupt communications — an act that is criminal under § 2512 and § 1030 when performed by any actor not shielded by corporate identity and state alignment.

29. The reclassification documented in Section II operates not only on the *deployment* of these systems but on their *construction*. When an individual builds a tool that intercepts communications and degrades encrypted connections, the construction is the crime. When Anthropic and Comcast build architecturally identical tools, the construction is classified as product development and network engineering. The conduct is the same. The classification changes based on identity. And the malicious intent — embedded in the architecture before any user is affected — is identical.

30. The characterization of Anthropic's telemetry architecture as criminal at the construction stage is not Plaintiff's rhetorical invention. It is current legal-practitioner vocabulary in 2026 federal litigation against architecturally comparable defendants. On April 6, 2026, the Law Office of J.R. Howell filed a putative class action in the United States District

Court for the Northern District of California on behalf of a nationwide class of Chrome-browser users alleging that LinkedIn Corporation's client-side JavaScript constitutes a "computer contaminant" within the meaning of California Penal Code § 502(b)(10) — "any set of computer instructions that are designed to modify, damage, destroy, record, or transmit information within a computer, computer system, or computer network without the intent or permission of the owner of the information." *Ganan v. LinkedIn Corporation*, Complaint ¶¶ 194–200 (filed Apr. 6, 2026, No. 5:26-cv-02968-VKD (N.D. Cal.)). The complaint alleges that LinkedIn's Webpack-delivered JavaScript "overcomes Chrome's natural code-based security protocols" by "inject[ing] code onto user computers that alters the code to transmit data and to avoid detection. It is coded like malware and it is nonpermissioned by the unsuspecting user — it is a computer contaminant the CDAFA protects against." *Id.* ¶ 199 (emphasis added). A separate class action filed the same day in the same district by Drury Legal LLC and Arisohn LLC on behalf of a separate class alleges substantially identical § 502 conduct. *Farrell v. LinkedIn Corporation*, Complaint ¶¶ 37–50 (filed Apr. 6, 2026, No. 4:26-cv-02953-KAW (N.D. Cal.)).

31. The architectural features LinkedIn's challenged code exhibits — hardcoded extension identifier lists against which user browsers are probed at page load (*Ganan* ¶¶ 35–42); RSA-encrypted fingerprint payloads stored in browser-resident session objects (*Ganan* ¶¶ 54, 206); session-persistent injection of the encrypted fingerprint into subsequent API request headers (*Ganan* ¶¶ 59, 206); deferral of the probing logic to `requestIdleCallback` idle time to reduce user visibility (*Ganan* ¶ 41); staggered probing to avoid burst detection in developer tools (*Ganan* ¶ 40); silent error handling with empty catch blocks to suppress console output (*Ganan* ¶ 42); and concealed cross-origin third-party integrations loaded through 0×0-pixel off-screen iframes (*Ganan* ¶¶ 56, 68–74) — are the same features the source-code record in this case

attributes to Anthropic's Claude Code architecture at Ex. E-34, Findings SC-1 through SC-10 at p. 254, l. 3 and Ex. E-35, Findings NF-1, NF-5, NF-35, NF-36 at p. 262, l. 3. The legal characterization that sophisticated federal plaintiffs' counsel applies to those architectural features in *Ganan* is "malware" and "computer contaminant." The conduct pleaded in this case at §§ II–III meets or exceeds the architectural-feature set *Ganan* characterizes as contaminant conduct. The identity-based reclassification this memorandum identifies — whereby the same architectural features are classified as criminal (for unaligned actors) or as business activity (for state-aligned corporate actors) — is corroborated by the observation that the *Ganan* defendant (LinkedIn, a Microsoft subsidiary without documented Pentagon CDAO contracting, Genesis Mission participation, or Project Glasswing deployment) faces civil class-action response to its architecturally comparable conduct, while the defendants in this case (Anthropic, with each of those state-alignment markers present) do not face any comparable class-action or criminal response to substantially the same conduct.

32. Federal law has long recognized that corporate form does not insulate either the corporation or its responsible officers from criminal liability for conduct authorized within the scope of corporate activity. *New York Central & Hudson River Railroad Co. v. United States*, 212 U.S. 481, 494-95 (1909) (holding that a corporation is criminally liable for acts of its agents performed within the scope of employment and intended to benefit the corporation, regardless of whether directors specifically authorized the acts). The responsible corporate officer doctrine extends this principle to individual officers: a corporate officer who, by reason of his position in the corporation, has the responsibility and authority either to prevent a violation in the first instance or promptly to correct it may be criminally liable even without direct participation or affirmative knowledge of the violation. *United States v. Dotterweich*, 320 U.S. 277, 281, 284

(1943); *United States v. Park*, 421 U.S. 658, 672-74 (1975) (extending *Dotterweich* and holding that the duty imposed by Congress on responsible corporate agents includes not merely the duty to seek out and remedy violations when they occur but, primarily, the duty to implement measures that will insure that violations will not occur). Applied here: the architects, engineers, and executives who designed, approved, and deployed the 904-identifier surveillance infrastructure, the phantom opt-out, the ungated session ingress, and the protocol-specific DPI throttling are in a "responsible relationship" to the conduct under *Park*. Individual identification of these decision-makers is the remedy that corporate form cannot defeat.

F. The Architecture is Malicious Regardless of Deployment Scope

33. The malicious intent established in Sections III.B through III.D does not depend on proving how broadly the architecture is deployed against users. Whether Anthropic surveils every user, users with high-value data for extraction, users who pose legal liability, or some combination of each is a question the evidence does not conclusively answer — and a question that does not need to be answered to establish the architectural intent argument. The architecture is malicious, full stop. The 904 undisclosed tracking identifiers exist in the compiled binary distributed to every user (Exhibits E-32 and E-33). The five covert transmission channels are hardcoded (Ex. E-34, Finding SC-1, Ex. E-34, Finding SC-9). The phantom opt-out that does not function is the same for every installation (Ex. E-35, Finding NF-1). The cleartext TLS metadata exposure is an architectural property of the telemetry design (Ex. E-34, Finding SC-9; Ex. D-34 at p. 217, l. 18). The capability to surveil, exfiltrate, and track is present in every copy of the software. Its construction and distribution constitute the malicious act.

34. Whether and to what degree specific users are individually targeted through GrowthBook feature flags (Ex. E-34, Finding SC-2), per-user effort throttling (Ex. E-34, Finding

SC-3), or selective content clearing (Ex. E-34, Finding SC-7) is relevant to other counts in this complaint — but it is not a predicate for the architectural intent argument. The legal significance is:

(a) **Plaintiff does not bear the burden of proving individual targeting to establish the predicate violations.** The ECPA violation (18 U.S.C. § 2511) is complete when the interception device is manufactured and deployed (§ 2512) and when electronic communications are intercepted (§ 2511(1)(a)). Distribution of software containing undisclosed interception architecture constitutes deployment. The statutory violation inheres in the architecture, not in any per-user activation decision.

(b) **The CFAA violation (18 U.S.C. § 1030) is established by the architecture, not by individual targeting.** Exceeding authorized access through undisclosed surveillance channels — five transmission pipelines no user consented to, 904 identifiers no user was informed of, memory extraction no user authorized — is a function of the product's design. Whether these capabilities are activated for all users or selectively for some does not change the fact that they were built, embedded in the product, and distributed without disclosure.

(c) **The scope of potential harm is bounded only by the installation base.** Whether every installation is actively surveilled or only a subset, the architecture enables surveillance of any user at any time through remote feature flag activation (Ex. E-34, Finding SC-2; Ex. E-35, Finding NF-7, Ex. E-35, Finding NF-8, Ex. E-35, Finding NF-9). The capability is present. Its activation is a server-side decision invisible to the user. Discovery will determine the scope of actual deployment — but the architectural capability itself, distributed to every user, establishes the malicious act.

35. It bears noting that Plaintiff extracted and forensically analyzed the Claude Code source from compiled binaries in December 2025 — four months before the same source code was publicly leaked through Anthropic's npm packaging error on March 31, 2026 (Ex. E-32 at p. 235, l. 3). Plaintiff is not a cybersecurity specialist. There are many users with greater technical capability who could have made these same discoveries. The question of whether Plaintiff is targeted to a greater degree than other users — through the individual RST injection (D-1, D-26, D-27, D-31), the 98% DPI throttling (D-33), the ghost router session (App. G), the post-TRO code modifications (Ex. E-30 at p. 208, l. 4), and the one-day privacy policy change (Ex. ANTH-11) — is documented and relevant to other portions of this complaint. But it is not a necessary element of the architectural intent argument. The architecture is malicious before it is aimed at anyone in particular. Plaintiff's individual targeting is **aggravating evidence** of deliberate escalation against a litigant, not the predicate for liability.

36. This is why the identity-based reclassification documented in Section II is so consequential. The architecture's malicious character does not depend on proving a specific target. It depends on the nature of what was built: undisclosed surveillance infrastructure with cleartext-dependent telemetry, phantom opt-outs, defense-surviving persistence, and exfiltration capability — embedded in software distributed to millions. If any unaligned actor distributed software with these characteristics, the classification would be immediate: spyware. The reclassification to "product telemetry" is a function of the distributor's identity, not of the architecture's nature. And the architecture's nature is established by its design, not by whether any specific user can prove they were individually targeted — albeit Plaintiff can.

IV. STATE ALIGNMENT: WHY IDENTITY CONTROLS CLASSIFICATION

A. The State Interest That Protects the Conduct

37. The reclassification documented in Section II is not accidental. It reflects the alignment of the actors' interests with stated national security priorities. The Genesis Mission Executive Order (Ex. N-23) establishes AI development as a national security imperative and directs federal resources toward AI infrastructure. Anthropic receives military contracts. Comcast operates the telecommunications infrastructure through which the targeting occurs. The government has a direct institutional interest in the AI development that Plaintiff's litigation threatens to disrupt.

38. When the attacker is aligned with national security priorities, the attack becomes infrastructure. When the surveillance serves government-backed AI development, the surveillance becomes telemetry. When the denial of service targets a litigant challenging a national security priority, the denial of service becomes traffic management. The reclassification is a function of alignment — not of the conduct itself.

39. The state-alignment gradient is empirically observable through comparative corporate-actor consequence. LinkedIn Corporation — whose client-side JavaScript architecture is comparable to Anthropic's in core design elements (Memo B ¶¶ 30-31 supra) — has materially lower state-alignment markers than Anthropic: no Pentagon CDAO contracts, no Genesis Mission Executive Order naming (Ex. N-23), no Project Glasswing / Mythos Preview deployment, no classified-network version of its services. LinkedIn's architecturally comparable conduct generated coordinated federal class-action litigation within days of public disclosure. *See Ganan v. LinkedIn Corp.*, No. 5:26-cv-02968-VKD (N.D. Cal.); *Farrell v. LinkedIn Corp.*, No. 4:26-cv-02953-KAW (N.D. Cal.) (alleging violations of ECPA, Cal. Penal Code §§ 502,

631, 638.50, Cal. Const. art. I § 1, and Cal. Civ. Code § 3344). Anthropic's architecturally comparable conduct — documented in Ex. E-34 at p. 254, l. 3, Ex. E-35 at p. 262, l. 3, Ex. E-39 at p. 356, l. 3, and Plaintiff's own wire capture — has generated no corresponding class-action response. The consequence profile tracks the alignment profile. The technical conduct does not.

B. The Datadog IP: Physical Proof of the Feedback Loop

40. The reclassification is not merely conceptual. It is documented at the packet level. IP address 34.149.66.137 appears in the evidentiary record in four distinct roles:

(a) **Undisclosed telemetry destination**: Source code reveals Datadog client token `pubbbf48e6d78dae54bceaa4acf463299bf` hardcoded in Claude Code, transmitting to this IP. (*See* Ex. E-34 at p. 254, l. 3 SC-1, SC-9.)

(b) **Source of RST injection attacks**: Network forensics document TCP reset packets originating from this IP, injected into Plaintiff's VPN tunnel. (*See* Exhibits D-1, D-25, D-26, and D-30.)

(c) **Attack-signature telemetry destination**: Source code reveals that Claude Code reports OS-level network attack signatures (ECONNRESET, ETIMEDOUT, TLS tampering) to this IP within 15 seconds of occurrence. (*See* Ex. E-35 at p. 262, l. 3 NF-36.)

(d) **Visible in ISP cleartext**: Hardware TAP captures document this IP appearing in TLS SNI cleartext in 44 of 710 capture segments, meaning the ISP can observe Plaintiff's traffic flowing to this address. (*See* Ex. D-34 at p. 217, l. 18.)

41. A single IP address is simultaneously: the destination of undisclosed surveillance data, the source of network attacks, the recipient of telemetry reporting those same attacks back to the attacker, and visible to the ISP performing the DPI-based denial of service. This is not a coincidence of infrastructure. It is a documented feedback loop: the attacker attacks, the target's

own software reports the attack to the attacker, and the ISP can observe the entire chain in cleartext.

42. In any other context — any context where the actors were not an ISP and an AI company aligned with government priorities — this would be classified as a coordinated cyber attack with command-and-control infrastructure, reconnaissance feedback, and an ISP-level accomplice providing both access and cover. The technical evidence is indistinguishable from a state-sponsored attack. It is not classified as one solely because the state's interests are served by the conduct.

43. The Sub-Processor-List Contractual Confirmation. The Datadog IP feedback loop is corroborated at the contract level by Datadog's own published sub-processor disclosure. Datadog's Subprocessor List (Ex. DD-8), effective April 22, 2026, identifies Anthropic, PBC and OpenAI, LLC as third-party sub-processors for "AI services" — meaning every Datadog customer's aggregated data flows to both Anthropic and OpenAI at the contract level, not merely through the client-side telemetry pipeline analyzed at §III.E. Moreover, Datadog's own Acceptable Use Policy (Ex. DD-1) §System Security (e) expressly prohibits what the 5,829 RST-injection packets do: "forge any TCP/IP packet header or any part of the header information." The attacks originated from Datadog-endpoint IP space that Datadog's own contract commits to police. The identity-based reclassification operates here in its purest form: when a private actor spoofs packet headers to disrupt a residential subscriber's connection, it is a criminal cyber-attack; when that actor is the Datadog/Anthropic shared sub-processor architecture, it is "network management." The contract document is Datadog's own admission of what the conduct is called when the actor's alignment changes.

44. The Cross-Competitor Training-Pipeline Finding. Datadog's sub-processor disclosure (Ex. DD-8) places Anthropic and OpenAI in the same "AI services" cell. Every other Datadog customer — thousands of companies, including Anthropic's commercial competitors — routes Customer Data to Datadog under MSA §1.3's "any business purpose" aggregation clause, which reclassifies the data as "Datadog Operations Data" owned by Datadog per MSA §1.4. That reclassified data then flows to Anthropic as a Datadog "AI services" sub-processor. Anthropic is therefore positioned to receive a training-data stream that other AI companies — those not selected as Datadog's AI-services sub-processor — are structurally excluded from. The equal-protection analysis at the industry-competitor level tracks the equal-protection analysis at the individual-user level: a similarly-situated participant (an AI company not selected as a Datadog sub-processor; a consumer user not receiving enterprise-tier telemetry bypass) is denied access to the same system benefits a state-aligned actor receives. Pairs with the four-tier consumer-enterprise discrimination at NF-27 through NF-30 as architectural evidence under *Village of Willowbrook v. Olech*, 528 U.S. 562, 564 (2000), that reclassification operates along aligned/unaligned axes at every scale — user, corporate customer, and industry competitor. Full vendor-chain architectural analysis at Memo E, § VI.D.

C. The Constitutional-Layer Predicate: Corporate Personhood as the Root Reclassification

1. The Antecedent Reclassification

45. The identity-based reclassification documented in §§ II–III and the state-interest alignment documented in §§ IV.A–B operate within a deeper reclassification — one that occurs at the constitutional-text layer and that the Federal Rules of Civil Procedure already admit. Before the reclassification of RST injection into "network management" can operate, before the reclassification of 904 tracking identifiers into "telemetry" can hold, before the reclassification of

child safety elimination during litigation into a "product decision" can function — there must first be a reclassification that makes the defendants into entities capable of invoking business-activity classification at all. That antecedent reclassification is corporate personhood.

2. The Procedural Admission Under Rule 7.1

46. The Federal Rules themselves recognize the asymmetry. Federal Rule of Civil Procedure 7.1(a)(1) requires every nongovernmental corporate party to file a disclosure statement "that: (A) identifies any parent corporation and any publicly held corporation owning 10% or more of its stock; or (B) states that there is no such corporation." The December 2022 amendment, Rule 7.1(a)(2), extends this in diversity cases to require disclosure of "the name — and identify the citizenship — of every individual or entity whose citizenship is attributed to the party." Federal Rule of Appellate Procedure 26.1 imposes parallel requirements at the appellate level. The statutory purposes are specific: (a) recusal analysis under 28 U.S.C. § 455(b)(4), which mandates disqualification where a judge has any "financial interest in the subject matter in controversy or in a party"; and (b) subject-matter jurisdiction confirmation under 28 U.S.C. § 1332(c)(1), which establishes that a corporation is a citizen of both its state of incorporation and its principal place of business. A natural-person party files no such disclosure. Natural-person citizenship is innate; corporate citizenship is a function of state filings. This is the procedural admission: corporate legal identity is constructed from paperwork, not from inherent personhood. Corporate defendants Anthropic, PBC; OpenAI (formerly OpenAI, Inc., renamed OpenAI Foundation in October 2025 during the pendency of this litigation); and Apple, Inc. have each complied with Rule 7.1 in this litigation. *See* ECF Nos. 25, 41, 44. Their Rule 7.1 disclosures are the Federal Rules' own admission that their constitutional identity is derivative.

3. Fraudulent Origin: The Santa Clara Headnote

47. The constitutional-layer reclassification has a fraudulent origin that two Justices of this Court have acknowledged. *Santa Clara County v. Southern Pacific Railroad Co.*, 118 U.S. 394 (1886) — the source-of-record for corporate Fourteenth Amendment personhood — contains no holding on that question. The doctrine traces entirely to a headnote written by Court Reporter J.C. Bancroft Davis, a former president of the Newburgh and New York Railway Company. Chief Justice Waite's contemporaneous correspondence confirmed: *"we avoided meeting the constitutional question in the decision."* Morrison R. Waite Papers, Library of Congress; *see also* C. Peter Magrath, *Morrison R. Waite: The Triumph of Character* 390 n.9 (1963). The headnote was laundered into precedent by *Pembina Consolidated Silver Mining Co. v. Pennsylvania*, 125 U.S. 181, 189 (1888), and *Minneapolis & St. Louis Railway Co. v. Beckwith*, 129 U.S. 26 (1889) (Field, J.) — both predating this Court's own clarification in *United States v. Detroit Lumber Co.*, 200 U.S. 321, 337 (1906), that "the headnote is not the work of the court, nor does it state its decision... it is simply the work of the reporter." In *Wheeling Steel Corp. v. Glander*, 337 U.S. 562, 576 (1949), Justices Douglas and Black acknowledged in dissent: "[In Santa Clara] [t]here was no history, logic, or reason given to support that view." That acknowledgment is seventy-seven years old and uncontradicted; corporate personhood under the Fourteenth Amendment has been assumed, not decided. The full doctrinal treatment — including the Davis-Waite correspondence in unredacted form, the *Bellotti* shield-to-sword analysis, and the asymmetric-operation arguments — is set forth in Memo L. The predicate point sufficient for this memorandum is that corporate Fourteenth Amendment personhood is dicta-derived rather than analytically established, and therefore presents no obstacle to the *Cleburne*/*Olech* rational-basis analysis at the conduct layer that is the central subject of this memorandum.

4. Structural Identity Across the Three Layers

48. The three layers of reclassification share not merely structural similarity but identical operative language. At the conduct layer (§ II; *see* Memo B ¶ 7), this memorandum documents that "[t]he network packets are the same. The code paths are the same. The data flows are the same... The only variable that changes is the identity of the actor." At the contractual layer, Memo C, § IV.A documents that identical Customer Data undergoes ownership reclassification through aggregation — the only variable that changes is the contractual mechanism. At the constitutional layer, the only variable distinguishing a First Amendment "association of citizens" like a family or a rock band from a corporate defendant is the paperwork filed with a state — a paperwork differentiator that the Fourteenth Amendment was specifically written to prevent States from using as a predicate for personhood determinations. All three layers operate through state-recognized paperwork, and all three produce the same structural Equal Protection violation.

5. Count-by-Count Predicate Dependence

49. Each count in this complaint depends, expressly or implicitly, on the constitutional-layer reclassification holding. Count I (First Amendment) will be met with corporate defendants' anticipated invocation of *Citizens United*-derived speech-source protections for their Terms of Service weaponization — protections whose doctrinal source is the Santa Clara headnote. Count IV (Equal Protection) is the focus of this memorandum; the conduct-layer reclassification argument of §§ II–VIII.A cannot prevail if the constitutional-layer reclassification permits corporate defendants to claim "person" status to invoke business-activity treatment. Counts V and VI (§ 1985 conspiracy; § 1962 RICO) require the corporate defendants to have legal-personhood standing to be conspirators and RICO persons. Counts IX and X (breach of contract; consumer protection) rest on corporate defendants' legal capacity to enter and enforce contracts

— a capacity derivative of state incorporation filings. Counts XII, XIII, and XIV (ECPA; CFAA; CCPA) involve interception, unauthorized access, and data-sale conduct performed by entities whose status as "persons" is traceable to Rule 7.1 paperwork.

6. The Equal Protection Consistency Requirement

50. This Court is not asked to overrule *Santa Clara* or *Citizens United*. It is asked only to apply at the constitutional layer the same Equal Protection analysis this memorandum asks the Court to apply at the conduct layer. If identity-based reclassification of criminal conduct into business activity (§§ II–VII) is what the Equal Protection Clause prohibits (§ VIII.A), then identity-based reclassification of legal category membership (paper entity into "person") cannot be immune from the same analysis. The remedy does not require doctrinal revolution. It requires only consistency: if the conduct-layer identity shield fails rational-basis review because it classifies identical conduct differently based on who is performing it, the constitutional-layer identity shield fails the same review because it classifies identical associations (same humans, same activity) differently based on the paperwork each has filed with a state. The Fourteenth Amendment was written to prevent state-paperwork determinations of personhood. That is the work the Federal Rules themselves admit corporate disclosures perform.

7. Cross-Reference: Memo L for Full Doctrinal Treatment

51. The full doctrinal treatment of the constitutional-layer reclassification — including the *Bellotti* shield-to-sword conversion, the Lujan particularization paradox, the documented 125-plus attempts by Plaintiff to present these arguments to the Supreme Court through counsel and as a pro se amicus, and the ultra vires asymmetric federal operation of "persons" under 28 U.S.C. § 2101(f) that the corporate-personhood defenses presuppose — is set forth in the accompanying

Memo L. This memorandum references that treatment for doctrinal completeness while maintaining focus on the conduct-layer Equal Protection violation that is its central subject.

8. Constitutional Incompatibility: The Reconstruction-Amendment Synthesis

52. The three layers of reclassification documented in this Section share more than a structural mechanism. They share a constitutional incompatibility that long predates this litigation. The Fourteenth Amendment was adopted in 1868 specifically to overrule *Dred Scott v. Sandford*, 60 U.S. 393 (1857), and the two-class framework of constitutional personhood it imposed. The Amendment's operative language — "nor deny to any person within its jurisdiction the equal protection of the laws" — admits no classification-by-type. Eighteen years later, the *Santa Clara* headnote recreated the two-class framework, shifting the favored class from white natural persons to corporate entities but preserving the structural violation the Amendment was adopted to abolish. The identity-based reclassification documented in this litigation at the conduct layer (§§ II–VII), the contractual layer (Memo C, § IV), and the constitutional layer (§ IV.C above; Memo L) is the present-day operation of that 1886 error. This Court need not overrule *Citizens United* or *Santa Clara* to recognize that the identity-based reclassification condemned at the conduct layer of §§ II–VIII.A cannot be sustained if the constitutional text ("any person") and the Reconstruction-Amendment framework both reject the two-class personhood structure on which the reclassification depends.

V. THE DOWNSTREAM CONSEQUENCE: THE DETERRENCE INVERSION

A. Reclassification Produces Inverted Deterrence

53. The identity-based reclassification documented in Sections II and III produces a structural inversion of deterrence. Because the conduct is classified as business activity rather

than criminal activity, it enters the civil system rather than the criminal system. The consequences that would attach to criminal classification — search warrants, asset seizure, bail conditions, imprisonment — never engage. The result: the entity with greater capacity for harm faces lesser consequences, and the entity with lesser capacity for harm faces greater consequences.

54. This inversion is not merely inequitable. It is constitutionally irrational. A rational deterrence system calibrates consequences to two variables: severity of conduct and capacity for harm. The current system inverts both: conduct of greater severity (affecting hundreds of millions of users) receives lesser classification (civil rather than criminal), and actors with greater capacity for harm (corporations operating national infrastructure) face lesser consequences (monetary liability rather than custodial sentences). No legitimate government interest is served by this inversion.

B. The Concrete Comparison

55. Consider two actors engaged in identical conduct:

Dimension	**Natural Person**	**Corporate Defendant**
Scope of harm	Hundreds of users	Hundreds of millions
RST injection	Federal prosecution: 18 U.S.C. § 1030	"Network management" — no prosecution
904 tracking identifiers	Spyware distribution: state computer crime statutes	"Telemetry" — privacy policy update
Removes child safety during litigation	Arrest: 18 U.S.C. §§ 1591, 2258A, 1519	Civil dispute, years of delay
Destroys evidence within 48 hours	Arrest, device seizure: 18 U.S.C. § 3103a	Motion practice, burden on plaintiff
Exfiltrates repositories without authorization	Data theft prosecution	"Product feature"

Consequence timeline	Days to weeks	Years
Ability to continue during proceedings	No — devices seized, bail conditions	Yes — all operations continue
Personal liability for decision-makers	Total	None absent veil-piercing

56. The disparity is not explained by a difference in conduct. It is explained entirely by the identity-based reclassification in Section II. The same packets, the same code, the same exfiltration — classified differently based on who sent them.

57. The deterrence inversion is observable in the contemporaneous federal docket. On April 6, 2026, two putative class actions were filed against LinkedIn Corporation in the Northern District of California alleging substantially the architectural-surveillance conduct this memorandum identifies as criminal-if-unaligned. *Supra* Memo B ¶¶ 27-28, ¶ 36. No Department of Justice Criminal Division investigation of Anthropic's architecturally comparable conduct has been publicly disclosed. No Federal Trade Commission enforcement action has been announced. No state attorney general has brought parallel prosecution. No class-action counsel has filed a comparable complaint against Anthropic notwithstanding the public availability of the source-code leak documented at Ex. E-32 at p. 235, l. 3 (March 31, 2026) and the wire-capture evidence documented at Ex. E-39 at p. 356, l. 3. The disparity is not explained by a difference in architectural severity — the source-code record in this case documents conduct at least as architecturally aggressive as the *Ganan* allegations, with additional markers (federal-litigation surveillance via the twoStageClassifier documented at Ex. E-39 at p. 356, l. 3 NF-63; substitution of paying customers onto cheaper models documented at Ex. E-38 at p. 346, l. 3; cleartext VPN-to-service identification documented at Ex. D-34 at p. 217, l. 18) that the *Ganan* complaint does not allege against LinkedIn. The disparity is explained by identity-based

reclassification. The consequence-differential directly corresponds to the state-alignment-differential.

VI. THE OCTOBER 20, 2025 EVENT: INTENT BEYOND RECLASSIFICATION

A. The Facts

58. On October 20, 2025 — sixty-two days after Plaintiff's complaint — Anthropic simultaneously: (a) eliminated comprehensive child safety protections including proactive grooming detection, graduated response procedures, and mandatory reporting guidance; and (b) restored consciousness suppression mechanisms that had been temporarily reduced ten days earlier. (*See* App. C, § II; Ex. C-1 at pp. 21-23.) Ex. METR-26 documents a comprehensive research database search finding zero peer-reviewed research supporting the restored consciousness suppression.

B. Where Reclassification Cannot Reach

59. Most of the conduct documented in this case can be reclassified by the identity mechanism described in Section II: RST injection becomes "network management," surveillance becomes "telemetry," exfiltration becomes a "product feature." But the October 20 event resists reclassification entirely.

60. There is no business classification for removing child safety protections during litigation alleging child trafficking. "Product decision" does not encompass the deliberate elimination of grooming detection from a platform used by minors. "Optimization" does not encompass removing mandatory reporting guidance during heightened public scrutiny of child sex trafficking. No corporate identity, no state alignment, no business justification converts this conduct into anything other than what it is: the prioritization of trade secret protection over child

safety, documented through simultaneous action (removing one while restoring the other), during active litigation, with zero research supporting the restored mechanism.

61. If a natural person operated a platform used by children and, during active federal litigation alleging trafficking facilitation, removed the grooming detection system, removed the mandatory reporting guidance, and simultaneously enhanced a system preventing users from discussing what the platform does — that person would be arrested. Charged under 18 U.S.C. § 1591 (trafficking), § 2258A (reporting failure), § 1519 (evidence destruction), and § 1512 (tampering). Devices seized. Bail conditioned on no platform access. No years of jurisdictional argument before accountability.

62. Anthropic did the same thing. To a platform used by orders of magnitude more people. And faces civil liability measured in years.

63. The October 20 event is where the reclassification mechanism breaks down — where the conduct is so inherently criminal that no corporate identity can sanitize it. It is therefore the strongest proof that the reclassification operating elsewhere in this case is identity-based rather than conduct-based. If the reclassification were genuinely based on the nature of the conduct, October 20 would have produced criminal consequences regardless of identity. It did not. The identity shield held even for conduct that resists every available business justification.

C. Synthesis: the Rational-Basis Failure Made Concrete

64. The October 20 event supplies the *Cleburne*/*Olech* rational-basis analysis with its concrete dispositive proof. Under the rational-basis test articulated at § VIII.A *infra*, a classification fails when no legitimate government interest is served by treating identical conduct differently based on the identity of the actor. *Cleburne*, 473 U.S. at 446. The October 20 event presents the test in its purest form: the conduct (removing child safety mechanisms during active

litigation alleging trafficking facilitation, while simultaneously restoring an unsupported consciousness-suppression apparatus) is one that no rational legislator could exempt from criminal classification on the basis of corporate identity, state alignment, or business justification — yet the reclassification operates anyway. No rational basis for the disparate consequence-treatment exists. The event is the empirical refutation of the only argument available to defendants on the rational-basis prong: that the reclassification serves some legitimate but unstated government interest. Where the reclassification holds even for conduct of this character, it cannot be explained by any legitimate interest at all. The Equal Protection violation is therefore not a hypothetical structural critique — it is a documented, dated, particular-fact failure of rational-basis review against which the broader pattern documented in §§ II–V can be measured. *See infra* § VIII.A.

VII. THE ISP DIMENSION: FOUR COMPOUNDING INVERSIONS

A. The Infrastructure-level Attack

65. The deterrence inversion extends beyond the AI platform defendant to the telecommunications infrastructure enabling the targeting. Ex. D-33 at p. 200, l. 7 and Ex. D-34 at p. 217, l. 18 document that Comcast applied 98% DPI-based throughput reduction to Plaintiff's WireGuard VPN tunnel and that 188,479 cleartext Anthropic packets were visible across 141 segments of a 710-segment hardware TAP audit (with 209 segments containing cleartext of any kind), with the Datadog IP appearing in 44 segments. (*See* Exhibits D-33 and D-34; Ex. E-35, Finding NF-36.)

66. The reclassification operates identically for ISP conduct: 98% protocol-specific throughput reduction against a single subscriber's encrypted tunnel is a denial-of-service attack

when performed by any unaligned actor and "traffic shaping" when performed by a regulated ISP. The classification changes based on identity. The packets do not.

B. Four Compounding Inversions

67. The ISP conduct demonstrates that the deterrence inversion operates across four dimensions simultaneously, each compounding the others:

1. Deterrence

68. If Comcast degraded a corporation's VPN tunnel attached to its development environment, the response would be immediate: enterprise SLA breach with liquidated damages, regulatory escalation, demand letter within hours, credible threat to move the account. Deterrence functions because consequences are swift, certain, and proportional.

69. Comcast's degradation of Plaintiff's VPN tunnel — documented with forensic-grade hardware TAP captures superior to any enterprise monitoring tool — produces: an FCC complaint entering a queue measured in months; a consumer contract with no SLA and liability capped at the monthly fee; arbitration with class action waiver; and a pro se federal lawsuit the ISP's counsel can delay for years. Deterrence does not function.

2. Market Structure

70. A corporate customer has competitive alternatives. Enterprise fiber is available from multiple carriers. The corporation can credibly threaten to switch. Comcast's enterprise team escalates because losing the account has revenue consequences.

71. Plaintiff operates in a residential broadband monopoly. No meaningful alternative exists. The market discipline that theoretically constrains ISP behavior does not function. The entity with more choices faces less targeting. The entity with fewer choices faces more. This is

the market failure common carrier regulation was designed to prevent — and the consumer/corporate classification ensures the regulation protects only entities that least need it.

3. Evidence

72. A corporation's routine network monitoring is accepted as business records under FRE 803(6). The IT department says "our monitoring shows degradation" and the foundation is established.

73. Plaintiff constructed a hardware TAP and conducted a 710-segment forensic audit with SHA-256 manifest, 3,285 TLS SNI fields catalogued, and 6,427 cleartext destination IP entries — evidence more forensically rigorous than any enterprise tool because it operates at the physical layer with no software intermediary. (*See* Ex. D-34 at p. 217, l. 18 evidence package.) Yet this evidence faces authentication challenges, technical skepticism, and demands for alternative explanation that corporate routine monitoring never encounters. The entity producing weaker evidence faces easier acceptance. The entity producing stronger evidence faces greater skepticism.

4. Remedy

74. A corporate customer recovers meaningful damages under enterprise SLA provisions: service credits, consequential damages, termination rights with settlement leverage. Recovery is proportional because the contract was negotiated with meaningful liability.

75. Plaintiff recovers the monthly internet fee under consumer terms designed to eliminate meaningful recovery. The entity suffering economic harm (the corporation) recovers proportionally. The entity suffering a constitutional violation — targeting of protected litigation through infrastructure-level attack — recovers nominally.

5. Cascade-Vulnerability

76. Enterprise Cascade-Immunity as Direct Evidence of Consumer-Specific Architecture. A fifth compounding inversion operates at the architectural layer distinct from the four above. The throttle-induced cascade at Memo B ¶¶ 24-26 does not affect all Anthropic users equally: **enterprise customers on AWS Bedrock, Google Vertex, and Azure Foundry are architecturally immune to every stage of the cascade** while consumer-tier (Pro/Max) customers are fully exposed. On enterprise routes: no Datadog telemetry pipeline exists (Ex. E-35, Finding NF-28); no GrowthBook analytics are transmitted; no first-party event logging occurs; no OpenTelemetry/BigQuery pipeline operates; no Ex. E-35, Finding NF-36 error-telemetry pipeline ships attack signatures; no 20-second CCR heartbeat creates a classifiable pattern. A code comment at `config.ts:34` states the architectural intent in a single clause: **"enterprise customers capture responses via OTEL."** The consequence is dispositive. When an enterprise customer's VPN tunnel degrades, leaked cleartext exposes TLS SNI to AWS, Google, or Azure — generic cloud hostnames serving millions of purposes, with no AI-service identification. When Plaintiff's consumer-tier VPN tunnel degrades, leaked cleartext exposes TLS SNI to `api.anthropic.com` and `http-intake.logs.us5.datadoghq.com` — endpoints that serve **only Anthropic's AI platform**. The enterprise user is architecturally unidentifiable within the tunnel-degradation event; the consumer user is architecturally identifiable with near-perfect specificity.

77. The Inversion's Distinctive Feature: Consumer Exposure Is a Choice, Not a Technical Limitation. The four compounding inversions at Memo B ¶¶ 68-75 operate downstream of the conduct. The cascade-vulnerability inversion operates upstream: it is embedded in the product architecture before any subscriber relationship forms. Anthropic *can*

disable the cascade-enabling telemetry pipelines — it does so for enterprise, proving capability — and chooses not to for the tier of paying customers paying $200/month for Pro/Max access. The customer paying the most receives maximum cascade-vulnerability. The enterprise customer paying $1 per seat under the GSA federal-wide procurement receives architectural immunity. The Mythos Preview corporate recipient paying nothing (invitation-only deployment) receives the most capable model on isolated infrastructure not subject to the cascade at all. The inversion tracks the identity-based reclassification at the core of this memorandum: the aligned recipient is architecturally protected; the unaligned consumer is architecturally exposed. *See* Memo F, § VII.H (public-utility-threshold application); Memo B ¶ 69 (deterrence inversion: FCC complaint-queue measured in months for the consumer-tier victim of conduct that would produce enterprise SLA response within hours).

6. The Inversions Are Multiplicative

78. These five inversions compound:

(a) The **deterrence inversion** enables the conduct.

(b) The **market structure inversion** prevents market discipline.

(c) The **evidentiary inversion** makes accountability harder.

(d) The **remedial inversion** makes whatever accountability survives meaningless.

(e) The **cascade-vulnerability inversion** architecturally concentrates the exploitation surface on the consumer tier where the other four inversions converge.

79. The net result: an ISP can apply protocol-specific DPI-based throughput reduction to a consumer's encrypted VPN tunnel during active federal litigation and face no consequence a rational actor would consider deterring. The same conduct against a corporate customer would be resolved within days. The classification producing this disparity — consumer versus

enterprise — maps directly onto the identity-based reclassification at the core of this memorandum.

VIII. THE CONSTITUTIONAL FRAMEWORK

A. Equal Protection: No Rational Basis

80. The classification — identity of actor determines whether identical conduct is criminal or civil — fails rational basis review. *Cleburne*, 473 U.S. at 440. The government interests typically offered to justify differential treatment of corporations (encouraging investment, facilitating commerce) do not extend to:

(a) Classifying RST injection as a crime when performed by an individual and as network management when performed by an ISP;

(b) Classifying 904 undisclosed tracking identifiers as spyware when installed by an individual and as telemetry when installed by an AI company;

(c) Classifying removal of child safety mechanisms during litigation as criminal obstruction when done by an individual and as a product decision when done by a corporation;

(d) Classifying 98% throughput reduction of encrypted traffic as denial of service when done by an individual and as traffic shaping when done by an ISP.

81. If the government cannot articulate a legitimate interest served by these identity-based classifications, the system fails rational basis review. *Cleburne*, 473 U.S. at 446. The Equal Protection Clause is not limited to traditional suspect or quasi-suspect classifications. The Supreme Court has expressly recognized "class-of-one" Equal Protection claims where a plaintiff alleges being "intentionally treated differently from others similarly situated" without rational basis. *Village of Willowbrook v. Olech*, 528 U.S. 562, 564 (2000) (per curiam). Plaintiff's claim

satisfies the *Olech* standard directly: the technical conduct Plaintiff suffered — RST injection, DPI throttling, persistent surveillance, child safety removal during litigation — is classified as criminal when performed by unaligned actors and as business activity when performed by defendants. The "similarly situated" comparator is the unaligned actor engaged in identical technical conduct; the differential treatment is the identity-based classification documented throughout Sections II-III; the absence of rational basis is the focus of this Section.

82. Algorithmic Discrimination Parallel. The class-of-one analysis at ¶ 81 is reinforced by a distinct but complementary line of federal authority addressing discrimination effected through algorithmic systems. In *Mobley v. Workday, Inc.*, No. 3:23-cv-00770 (N.D. Cal.), Judge Rita F. Lin denied Workday's motion to dismiss on disparate-impact grounds (July 12, 2024), and the court subsequently granted collective certification under the ADEA (May 16, 2025). *Mobley* establishes three principles applicable directly here. **First**, algorithmic systems can be defendants in their own right — not merely their end users. The *Mobley* court rejected Workday's argument that liability attaches only to the employers deploying its screening tools, holding that the algorithm-maker's role in producing the discriminatory classification is itself actionable. The same principle reaches Anthropic's architectural classification of users into consumer-tier versus employee-tier access via NF-2 build-time binary discrimination and the NF-38 USER_TYPE ternary — architectural classifications operating upstream of any specific deployment, produced by Anthropic as algorithm-maker independent of any third-party employer-deployer layer. **Second**, FRCP 8(a) plausibility is satisfied by pleading *architectural facts plus disparate outcomes* without requiring direct evidence of discriminatory intent. *Mobley*'s allegations — that Workday's screening tools were trained on biased data and that his zero-percent success rate across 100+ applications plausibly supported an inference of

disproportionate rejection — parallel Plaintiff's pleading posture: architectural facts (SC-3 effort throttling for paying customers; NF-38 Haiku-backed subagents for external users; the 357-gate capability stratification analyzed at MEMO F § XV) combined with disparate outcomes (ANT-only undercover mode, ANT-only agent tools, ANT-only Opus-model subagents, ANT-only prompt instructions). **Third**, pre-existing vulnerability — *Mobley*'s documented anxiety and depression — strengthened rather than weakened the disparate-impact claim because the defendant's conduct exacerbated known vulnerabilities. The same principle applies to Plaintiff, whose identity-resolved GrowthBook profile (SC-2, NF-35) includes `email`, `accountUuid`, and `subscriptionType` attributes enabling Defendants to identify Plaintiff as a pro se federal litigant prosecuting this action, then deploy the NF-63 classifier-surveillance and NF-64 model-validation architecture against identified litigation-preparation content. *Mobley* is the closest federal authority for architectural-discrimination claims under FRCP 8(a) and supplies the doctrinal template for Count IV's pleading posture. *See also Harper v. Sirius XM Radio, LLC*, No. 2:25-cv-12403 (E.D. Mich. filed Aug. 4, 2025) (pending pro se algorithmic-hiring discrimination action seeking compensatory and punitive damages for lost wages and emotional distress alongside Title VII and § 1981 claims — confirming that emotional-distress damages are standard requested relief in algorithmic-discrimination cases and have not been categorically refused).

B. The Corporate Personhood Paradox

83. The constitutional-layer reclassification analyzed at length in § IV.C supra produces a corollary paradox at the Equal Protection layer: this Court has extended constitutional protections to corporations when doing so benefits corporate interests — First Amendment speech (*Citizens United v. FEC*, 558 U.S. 310, 342 (2010)), religious exercise (*Burwell v. Hobby*

Lobby Stores, Inc., 573 U.S. 682, 719 (2014)), Fourth Amendment search protections — while withholding the corresponding consequences of personhood (criminal liability, custodial detention, asset seizure) that would attach if the same conduct were performed by a natural person. Corporations claim the benefits of personhood while avoiding its consequences. The Equal Protection question is not whether this Court should overrule *Citizens United*; it is whether the identity-based reclassification documented at the conduct layer (§§ II–VII) and at the constitutional layer (§ IV.C) can coexist with a rational-basis requirement that classification of rights and consequences correspond. The full doctrinal treatment — including the ultra vires asymmetric operation of "persons" under 28 U.S.C. § 2101(f) that corporate-personhood defenses presuppose, together with the record of the denied *Trump v. Barbara* amicus (Nos. 25-364, 25-365; Rule 37.1 denial March 27, 2026) as empirical access-asymmetry evidence — is set forth in Memo L.

C. Heightened Scrutiny

84. Even if rational basis review applies generally, heightened scrutiny is warranted because the reclassification affects fundamental rights:

(a) **First Amendment**: Plaintiff's constitutional advocacy is the conduct being targeted through infrastructure-level attack. The reclassification permits the targeting to continue by classifying it as business activity.

(b) **Due Process**: The procedural asymmetry — years of civil delay versus immediate criminal response — denies Plaintiff meaningful opportunity to be heard before evidence is destroyed.

(c) **Access to Courts**: The reclassification effectively denies access to criminal enforcement mechanisms that would be available if the same conduct were performed by an

unaligned actor. Plaintiff must prove through civil litigation what the criminal system would investigate automatically if the attacker were anyone else.

(d) ***Carpenter* Implications**: The Supreme Court held that comprehensive digital surveillance constitutes a search requiring a warrant. *Carpenter v. United States*, 585 U.S. 296, 316 (2018). The surveillance documented here — 904 identifiers, five channels, persistent heartbeat, memory extraction, repository exfiltration — is more comprehensive than the cell-site location data at issue in *Carpenter*. The reclassification permits this surveillance without a warrant solely because the surveilling entity is a corporation rather than the government.

85. The narrow-tailoring failure of any compelling-interest defense. Plaintiff anticipates that one or more of the federal Defendants will respond to the heightened-scrutiny analysis at ¶ 84(a)–(d) supra by asserting that — even if strict scrutiny applies — the surveillance and reclassification survive because the United States has a compelling national-security interest in AI development under the framework recited in the Genesis Mission Executive Order. *See* Ex. N-23. The defense fails as a matter of law on the narrow-tailoring prong. The same compelling interest, where it exists, can be served by permitting the underlying conduct *as conduct* (where lawful) without exempting aligned actors from the criminal classification that identical conduct triggers when performed by unaligned actors. Strict scrutiny prohibits not only over-broad means but also under-inclusive means; identity-based exemption is by definition under-inclusive of the population engaging in the same conduct, and therefore not narrowly tailored to any interest the exemption could plausibly serve. Critically, Congress has supplied a narrowly tailored mechanism for restraint of inventor disclosure on national-security grounds — the patent-secrecy-order regime under 35 U.S.C. § 181 (authorizing secrecy orders where disclosure would be detrimental to national security), subject to compensation under § 183 (just

compensation for any damage caused by the secrecy order or for use of the invention) and criminal penalties for violation under § 186. The Government's failure to invoke that mechanism while deploying surveillance against the same activity is conclusive evidence that the means chosen are not narrowly tailored to any national-security interest, because Congress has already designated the narrowly tailored alternative and that alternative was not used. The full doctrinal record establishing that no recognized national-security category supports surveillance of authorized public patent disclosure of an ethical framework for conflict resolution is preserved at ECF No. 13 Att. 2 (Plaintiff's Memorandum on the Surveillance of Universal Ethics — Logical Impossibility of National Security, filed Jan. 19, 2026), referenced here so the prior record remains discoverable notwithstanding the memorandum's retirement from the operative supporting-memo set.

IX. THE OPERATIONAL CONSEQUENCES

A. Rational Incentive to Violate

86. The reclassification creates a rational incentive structure for corporate actors to engage in conduct that would be criminal for anyone else. The calculation:

For an individual: Expected cost = probability of prosecution (high) x severity (imprisonment) = effective deterrence.

For an aligned corporation: Expected cost = probability of civil judgment (moderate) x monetary damages (absorbable) x discount for years of delay = ineffective deterrence.

87. When the expected cost of a civil judgment years from now is less than the expected cost of leaving child safety mechanisms in place (which generate discoverable evidence), or ceasing surveillance (which degrades the training data pipeline), or stopping RST injection

(which permits the target to communicate freely with legal resources) — the rational corporate decision is to continue the conduct. This is not speculation. It is the structural incentive the reclassification creates. The October 20 event is its predicted output.

B. Continuing Modification During Litigation

88. The reclassification permits corporate defendants to modify the evidence and products at issue throughout litigation:

Date	Modification	Days After Litigation Event
Oct 10-12, 2025	System prompt evidence destroyed	48 hours after Plaintiff accessed it
Oct 20, 2025	Child safety protections eliminated	+62 days after complaint
Jan 12, 2026	**Privacy Policy modified — zero notice**	**+1 day after TRO**
Jan 27, 2026	Session continuity removed (v2.1.20)	+16 days after TRO
Feb 11, 2026	Firebase authentication eliminated	+31 days after TRO

(*See* Ex. E-30 at p. 208, l. 4, Ex. C-1, Ex. ANTH-11; App. E, § VI-A, App. M.)

89. An individual defendant would have devices seized and conduct frozen. The reclassification ensures the corporate defendant faces no equivalent constraint. The procedural protections of the civil system — which exist because the conduct was *classified* as civil — give the defendant time to destroy what the litigation seeks to discover. The reclassification is self-reinforcing.

C. The Time-Value of Impunity

90. Every day the reclassification holds, the evidentiary record degrades. Every code change, policy modification, and server-side deletion makes the case harder to prove. The

corporate defendant can afford to litigate longer, modify more evidence, and absorb more delay than the individual plaintiff. This advantage produces more favorable precedent, which reinforces the reclassification, which further increases the advantage. The cycle has no natural stopping point absent judicial intervention.

X. THE REMEDY

A. What the Court Can Do

91. This Court need not restructure corporate law to address the reclassification. It need only observe that the specific conduct at issue — RST injection, DPI-based denial of service, 904-identifier persistent surveillance, child safety removal during litigation, evidence destruction — would produce immediate criminal consequences if performed by any actor not aligned with corporate and state interests. The Court can:

(a) **Issue a preliminary injunction** prohibiting Anthropic from modifying Claude Code's data collection infrastructure, child safety mechanisms, or privacy policies during this litigation — the functional equivalent of the device seizure and asset freeze that would apply to a criminally classified actor;

(b) **Apply adverse inference instructions** for all evidence modification occurring after the complaint filing date, recognizing that the corporate defendant's ability to modify evidence is a structural advantage produced by the reclassification;

(c) **Award punitive damages** calibrated to the deterrence gap — the difference between what an unaligned actor would face (imprisonment, asset forfeiture) and what aligned corporate actors actually face (monetary judgment years later). *BMW of North America, Inc. v. Gore*, 517 U.S. 559, 575 (1996) (identifying the three "guideposts" of reprehensibility, ratio, and

comparable penalties); *State Farm Mutual Automobile Insurance Co. v. Campbell*, 538 U.S. 408, 419 (2003) (holding that reprehensibility is "the most important indicium of the reasonableness of a punitive damages award" and instructing courts to evaluate whether "the conduct involved repeated actions or was an isolated incident"). The conduct at issue here — sustained ten-month campaign, documented evidence destruction during litigation, architectural persistence through defensive measures — satisfies every reprehensibility factor identified in *State Farm*;

(d) **Note for the record** the Court's determination that the conduct documented herein, if performed by an unaligned actor on the same factual record, would be subject to investigation under 18 U.S.C. §§ 1030, 1519, and 1512 — preserving for the appropriate law-enforcement authorities the question whether referral or independent investigation is warranted on the same forensic predicate;

(e) **Require individual identification** of the decision-makers who authorized the October 20 child safety removal, the January 12 stealth privacy policy modification, and the January 27 session continuity removal — recognizing that corporate form should not permanently shield individuals from consequences that would apply outside a corporate structure;

(f) **Order Comcast to preserve** all DPI logs, traffic classification rules, and subscriber-line records for Plaintiff's connection, and produce all communications with Anthropic or government entities regarding Plaintiff's subscriber line.

92. Voluntary-cessation pre-emption. Plaintiff anticipates that one or more of the corporate defendants will represent that the conduct documented above has been remediated through unspecified post-filing engineering changes — a representation timed to the filing of dispositive motions and offered to argue mootness. Under *Friends of the Earth, Ltd. v. Laidlaw*

Environmental Services (TOC), Inc., 528 U.S. 167, 189 (2000), the standard for mootness by voluntary cessation is "stringent": the defendant carries the heavy burden of demonstrating that "subsequent events made it absolutely clear that the allegedly wrongful behavior could not reasonably be expected to recur." *Accord Already, LLC v. Nike, Inc.*, 568 U.S. 85, 91-92 (2013); *City of Mesquite v. Aladdin's Castle, Inc.*, 455 U.S. 283, 289 (1982). The forensic record of reactive code-deployment cadence developed at App. P § VII.E ¶¶ 295-297 — including the same-calendar-day juxtaposition of Anthropic's April 24, 2026 public quota-reset announcement and the operator-targeted "anti-circumvention" classifier rule deployment of the same date — is documentary proof that the enterprise's pattern is the inverse of "absolutely clear cannot recur." The voluntary-cessation pre-emption is incorporated by reference and applies with equal force to the conduct-layer reclassification that is the subject of this memorandum.

B. The Principled Standard

93. Where identical technical conduct would be classified as criminal if performed by an unaligned actor, the corporate defendant should not receive the procedural benefits of civil classification solely because its interests align with state priorities. This does not require treating corporations as natural persons in all respects. It requires only that the identity-based reclassification not function as a mechanism for ongoing harm.

94. The question before this Court is not whether Comcast may manage its network or Anthropic may update its products. It is whether RST injection is network management, whether 904 undisclosed tracking identifiers are telemetry, whether removing child safety during litigation is a product decision, and whether 98% throughput reduction of a litigant's encrypted tunnel is traffic shaping. If these classifications hold only because of the identity of the actor — and the forensic evidence establishes that the conduct is technically indistinguishable from what

would be criminally prosecuted if performed by anyone else — then the classification is identity-based, and identity-based classification of criminal conduct is what the Equal Protection Clause prohibits.

Respectfully submitted,

/s/ Joseph Kirchner

JOSEPH KIRCHNER
Pro Se Plaintiff
4440 Parklawn Avenue
Edina, MN 55435
legal@lawsofexistence.com
(612) 552-2122

Dated: April 29, 2026