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 C: DATA LAUNDERING THROUGH THIRD-PARTY AGGREGATION — 18 U.S.C. § 1956 APPLIED TO PERSONAL DATA

TABLE OF CONTENTS

TABLE OF AUTHORITIES

CASES

Citation
Already, LLC v. Nike, Inc., 568 U.S. 85, 91-92 (2013)
Armendariz v. Foundation Health Psychcare Services, Inc., 24 Cal. 4th 83 (2000)

Boyle v. United States, 556 U.S. 938, 944-46 (2009)
City of Mesquite v. Aladdin's Castle, Inc., 455 U.S. 283, 289 (1982)
FTC v. Avast Limited, Docket No. C-4803 (FTC Feb. 22, 2024)
FTC v. GoodRx Holdings, Inc., FTC Docket No. C-4798 (FTC Feb. 1, 2023)
FTC v. Gravy Analytics, Inc. & Venntel, Inc., FTC Docket No. C-4814 (FTC Dec. 3, 2024)
FTC v. Kochava Inc., No. 2:22-cv-00377-BLW (D. Idaho filed Aug. 29, 2022)
FTC v. X-Mode Social, Inc. / Outlogic, LLC, FTC Docket No. C-4807 (FTC Jan. 9, 2024)
Friends of the Earth, Ltd. v. Laidlaw Environmental Services (TOC), Inc., 528 U.S. 167, 189 (2000)
H.J. Inc. v. Northwestern Bell Telephone Co., 492 U.S. 229, 239 (1989)
Interstate Circuit, Inc. v. United States, 306 U.S. 208, 226-27 (1939)
People of California v. Sephora USA, Inc., No. CGC-22-601380 (Cal. Super. Aug. 24, 2022)
Sedima, S.P.R.L. v. Imrex Co., 473 U.S. 479, 493 (1985)
Specht v. Netscape Communications Corp., 306 F.3d 17, 29-32 (2d Cir. 2002)
United States v. Faiella, 39 F. Supp. 3d 544 (S.D.N.Y. 2014)
United States v. Harmon, No. 19-cr-395 (D.D.C. 2021)
United States v. Ulbricht, 31 F. Supp. 3d 540 (S.D.N.Y. 2014)
Williams v. Walker-Thomas Furniture Co., 350 F.2d 445, 449 (D.C. Cir. 1965)

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)

FEDERAL STATUTES

Citation
15 U.S.C. § 45(a) (Federal Trade Commission Act § 5)
18 U.S.C. § 1030 (Computer Fraud and Abuse Act)
18 U.S.C. § 1341 (Mail Fraud)
18 U.S.C. § 1343 (Wire Fraud)

18 U.S.C. § 1956 (Laundering of Monetary Instruments)
18 U.S.C. §§ 1961-1968 (Racketeer Influenced and Corrupt Organizations Act)
18 U.S.C. § 2511 (Electronic Communications Privacy Act — Interception)

STATE STATUTES

Citation
Cal. Civ. Code § 1798.100(c) (California Privacy Rights Act — Data minimization)
Cal. Civ. Code § 1798.110 (California Privacy Rights Act — Right to know)
Cal. Civ. Code § 1798.120 (California Privacy Rights Act — Opt-out right)
Cal. Civ. Code § 1798.135(b) (California Privacy Rights Act — Opt-out signals)
Cal. Civ. Code § 1798.140(ad) (California Privacy Rights Act — Sale definition)
Cal. Civ. Code § 1798.140(ag) (California Privacy Rights Act — Service provider)
Cal. Civ. Code § 1798.140(ah) (California Privacy Rights Act — Sharing definition)
Cal. Civ. Code § 1798.140(m) (California Privacy Rights Act — Deidentified)
Cal. Civ. Code § 1798.140(v) (California Privacy Rights Act — Personal information)

REGULATORY AUTHORITIES

Citation
11 CCR § 7002(b)(5) (CPRA Consumer reasonable expectations)
11 CCR § 7002(d) (CPRA Data minimization — proportionality test)
11 CCR § 7050(a)(3) (CPRA Service provider permitted uses)
11 CCR § 7050(b) (CPRA Cross-context behavioral advertising)
11 CCR § 7051(a)(2) (CPRA Contract specificity requirement)
11 CCR § 7051(c) (CPRA Business due diligence)
16 C.F.R. § 318 (Health Breach Notification Rule)
CPPA Enforcement Advisory No. 2024-01 (April 2, 2024)
FinCEN Guidance on Virtual Currencies, FIN-2013-G001 (March 18, 2013)

FEDERAL RULES OF CIVIL PROCEDURE AND EVIDENCE

Citation
Fed. R. Civ. P. 8(a) (Plausibility Pleading Standard)
Fed. R. Civ. P. 26(b) (Discovery Scope and Proportionality)
Fed. R. Evid. 801(d)(2)(D) (Statement by Party-Opponent's Agent)

SECONDARY AUTHORITIES

Citation
Brennan Center for Justice, *Closing the Data Broker Loophole* (2024)
CRS Report, *Money Laundering: An Overview of 18 U.S.C. § 1956*, RL33315
CRS Report, *Virtual Currencies and Money Laundering*, R45664
Restatement (Third) of Agency §§ 8.01-8.02 (Fiduciary Duties — Loyalty, Information, Conflicts)
Tikkinen-Piri et al., *Data Laundering and Regulatory Arbitrage in Data Protection* (academic literature, 2024)
ZDNET, *Data Laundering: What Is It and Why Should You Care?* (2019)

I. INTRODUCTION

1. This memorandum establishes that Anthropic PBC's systematic transfer of user personal data to third-party recipients GrowthBook, Inc. and Datadog, Inc. — under contractual terms that reclassify restricted user data into unrestricted commercial assets owned by the recipients — constitutes data laundering analogous to, and potentially prosecutable under, the federal money laundering statute, 18 U.S.C. § 1956. The mechanism is structurally identical to financial money laundering: proceeds of unlawful activity (personal data collected through deceptive practices) are routed through intermediary transactions (third-party aggregation contracts) designed to conceal or disguise the nature and ownership of those proceeds (converting "Customer Data" subject to training prohibitions into "Datadog Operations Data" or GrowthBook "Aggregated Data" free of restrictions), and the laundered proceeds promote further unlawful activity (AI

products trained on laundered data are sold back to Anthropic, funding continued deceptive collection).

2. Documentary Foundation. This memorandum analyzes the contractual instruments governing the data transfers (Datadog Master Subscription Agreement, Datadog Data Processing Addendum with CCPA Appendix C, GrowthBook Customer Agreement, GrowthBook Data Protection Addendum with CCPA Schedule 4), the source code forensics establishing the scope of data transmitted (Ex. E-34 at p. 254, l. 3 findings SC-1, SC-2, SC-9, SC-10; Ex. E-35 at p. 262, l. 3 findings NF-1, NF-5, NF-9, NF-13, NF-19, NF-20, NF-22, NF-23, NF-25, NF-35, NF-36), the network forensic evidence establishing the transmission pathways (D-1, D-26, D-27, D-33, D-34), Anthropic's corporate policy documents establishing the deceptive representations (Ex. ANTH-1, Ex. ANTH-2, Ex. ANTH-7, Ex. ANTH-9, Ex. ANTH-11), and Datadog's own public statements confirming that customer telemetry data trains proprietary AI models sold commercially. All contracts archived at `3rd_party_contracts/`; research findings detailed in `THIRD_PARTY_RECIPIENT_RESEARCH_20260414.md`.

3. This memorandum should be read in conjunction with Memo E (Essential Purpose Failure), which establishes that the TOS ownership provisions are illusory; App. E (Reactive Policy Changes), which documents the complaint-triggered policy reversals implemented within days of the August 19, 2025 filing; Exhibits DD-4, DD-5, GB-2, and GB-6 (Third-Party Contract Evidence), which reproduce the critical contract clauses with exhibit-formatted authentication; `DATA_BROKERAGE_CONTRACT_ANALYSIS_20260414.md`, which provides the full contract analysis, CPRA service provider exception failure, data minimization framework, forensic evidence mapping, and 26 discovery demands; and `THIRD_PARTY_RECIPIENT_RESEARCH_20260414.md`, which documents Datadog's

commercial AI training pipeline — including the Toto foundation model (2.36 trillion time series points per arxiv:2505.14766; analyzed at ¶¶47 et seq.) and the Bits AI agent suite (which Datadog's own marketing describes as "learns from customers' data") — and GrowthBook's warehouse-native privacy claims contradicted by source code evidence.

4. The argument proceeds in eight parts: (1) the *Faiella* line of cases establishes §1956 is asset-agnostic — applying to any medium of value, not only traditional currency (§II); (2) personal data constitutes "proceeds" with documented market value derived from specified unlawful activity, i.e., wire fraud through deceptive data-collection practices (§§II.C, III); (3) the contractual aggregation mechanisms are "financial transactions" designed to conceal the restricted nature of the data, operating through a dual-contract architecture in which broad commercial grants in primary agreements are contradicted by narrow CCPA restrictions in parallel DPAs (§IV); (4) the resulting AI products constitute promotional laundering by funding continued deceptive data collection through a circular economy, with Anthropic receiving five distinct categories of material benefit from the Datadog pipeline (§V.A–C); (5) the architecture extends across Anthropic's full sub-processor network — GrowthBook, Sentry, Statsig, Segment, Honeycomb, OpenTelemetry/BigQuery — operating through two structurally distinct concealment mechanisms: internal dual-contract reclassification (Datadog pattern) and inter-vendor delegation (Sentry pattern), establishing that the laundering occurs in the gaps between contracts that no single vendor's contract can restrict (§V.C); (6) the multi-vendor architecture's empirical consequences are documented in Plaintiff's contemporaneous testimonial record (approximately 281 files, ECF Nos. 5-30, 5-33, 5-34), with architectural capacity proved by the findings in this memorandum corroborating victim-level observations and establishing the §1956(a)(1)(B)(i) concealment-intent element at the human-interaction layer (§V.C); (7) the

laundering operates through user-funded infrastructure, ISP co-liability at Comcast, and the closed-loop weaponization of laundered data in Anthropic's Mythos model deployed exclusively to ~40 Project Glasswing corporations (§§VI, IX, X); and (8) the scheme supports existing Counts V, VI, IX, X, XII, XIII, and XIV, with remedies reaching Anthropic, Datadog, GrowthBook, and the broader sub-processor network (§§XI, XII).

II. LEGAL FRAMEWORK: §1956 IS ASSET-AGNOSTIC

A. The Statutory Text

5. Section 1956(a)(1) prohibits conducting a "financial transaction" involving "the proceeds of specified unlawful activity" with either (A)(i) intent to promote the unlawful activity, or (B)(i) knowledge that the transaction is "designed in whole or in part **to conceal or disguise the nature, the location, the source, the ownership, or the control**" of the proceeds. The statute uses the term "**property involved in a financial transaction**" — not "money," not "currency," not "legal tender." A "financial transaction" includes any transaction "involving the movement of funds by wire or other means" or "involving one or more monetary instruments." § 1956(c)(4)–(5). The term "funds" is not defined. "Proceeds" means "any property derived from or obtained or retained, directly or indirectly, through some form of unlawful activity, including the gross receipts of such activity." § 1956(c)(9). The breadth of "any property" is deliberate. Penalty: fine of not more than $500,000 or twice the value of the property involved, whichever is greater, or imprisonment up to twenty years, or both. § 1956(a)(1).

B. Faiella and the Asset-Agnostic Principle

6. In *United States v. Faiella*, 39 F. Supp. 3d 544 (S.D.N.Y. 2014), the defendant operated an underground Bitcoin exchange on Silk Road and was charged with money laundering

conspiracy under §1956(h). The defendant moved to dismiss, arguing that Bitcoin does not qualify as "money" or "funds" under the statute.

7. The court rejected this argument, holding that Bitcoin transactions qualify as "financial transactions" under §1956 because they involve "the movement of funds." The court reasoned that "**funds**" includes "**money**," which refers to "**an object used to buy things.**" Because Bitcoin can be used to buy things, Bitcoin transactions involve "the movement of funds." The court's conclusion was unequivocal: "**one can money launder using Bitcoin.**"

8. The significance of *Faiella* is that §1956 does not require the laundered asset to be traditional currency. Bitcoin is not legal tender. It is not issued by any government. It has no intrinsic value. Its "currency" designation is colloquial, not legal. Yet the court held it qualifies as "funds" because it is used as a medium of exchange — **an object used to buy things**. The statute is asset-agnostic.

9. This principle was reinforced in *United States v. Ulbricht*, 31 F. Supp. 3d 540 (S.D.N.Y. 2014), where the court held that the Silk Road platform's facilitation of Bitcoin transactions constituted money laundering. In *United States v. Harmon*, No. 19-cr-395 (D.D.C. 2021), the D.C. District Court — the same court presiding over this case — applied the same reasoning to a Bitcoin mixing service, holding that the service's function of obscuring the source of Bitcoin proceeds constituted concealment laundering under §1956(a)(1)(B)(i).

10. FinCEN Guidance FIN-2013-G001 (March 18, 2013) further established that virtual currency exchangers are "money transmitters" subject to anti-money-laundering obligations, regardless of whether the virtual currency constitutes "legal tender." Congressional intent, as the *Faiella* court noted, was to close "**a gaping hole in the money laundering deterrence effort**"

— not to create a narrow statute that criminals could circumvent by choosing a non-traditional medium of exchange.

C. Personal Data as an Asset with Documented Market Value

11. If cryptocurrency — which is not legal tender, is not issued by any government, and has no intrinsic value — qualifies as "funds" under §1956 because it is "an object used to buy things," then personal data with documented market value satisfies the same test.

12. Personal data is commercially traded on established markets with documented per-record pricing:

Data Type	**Market Price per Record**
Raw device IDs (MAID only)	$0.001–$0.005
Enriched device profiles (ID + demographics + behavior)	$0.05–$0.50
Location-enriched (device ID + GPS trails)	$0.10–$2.00/month
Identity-resolved (linked to real identity, email, address)	**$0.50–$5.00**
Specialized segments (health, financial, political)	$0.25–$1.00+

13. The data Anthropic transmits to GrowthBook — including `email`, `deviceId`, `accountUuid`, and `subscriptionType` (Ex. E-34, Finding SC-2, Ex. E-35, Finding NF-35) — constitutes **identity-resolved** personal data, the highest-value category at $0.50–$5.00 per record. This data is transmitted as the explicit consideration in a commercial relationship: Anthropic receives feature-flagging services; GrowthBook receives an "unrestricted, non-exclusive right to use Customer Data, including for the purpose of training and improving GrowthBook's artificial intelligence algorithms" (Customer Agreement §4.2).

14. The data Anthropic transmits to Datadog — including behavioral events, error classifications, streaming patterns, and the 865-line event schema (SC-9, SC-10, NF-36) — is used by Datadog in its commercial AI training pipeline — including the Toto foundation model (2.36 trillion time series points per arxiv:2505.14766; analyzed at ¶¶47 et seq.) and the Bits AI agent suite (which Datadog's own marketing describes as "learns from customers' data"), which Datadog sells commercially to 31,400+ customers. Datadog's CPO stated: "Datadog is uniquely positioned to deliver value with AI as a platform that has **a wealth of clean, rich data** — we process **trillions of data points**." The data is Datadog's self-described competitive advantage — its commercial value is not hypothetical but is the basis of a $3.41 billion annual revenue business.

15. Personal data is "an object used to buy things" under the *Faiella* test: it is exchanged for analytics services, it trains AI products sold commercially, and it has documented per-record market pricing. It is used as consideration in commercial contracts just as cryptocurrency is used as consideration in transactions. If Bitcoin qualifies as "funds" for §1956 purposes, identity-resolved personal data with higher established market value qualifies equally.

III. THE "SPECIFIED UNLAWFUL ACTIVITY": DECEPTIVE DATA COLLECTION

A. Wire Fraud as the Predicate

16. Section 1956(c)(7)(A) defines "specified unlawful activity" (SUA) to include any offense constituting "a racketeering activity" under 18 U.S.C. §1961(1). Wire fraud under 18 U.S.C. §1343 is a RICO predicate offense under §1961(1)(B). Wire fraud requires: (1) a scheme to defraud, (2) use of wire communications in furtherance of the scheme, and (3) material misrepresentations or omissions.

17. Anthropic's data collection practices constitute wire fraud through the following material misrepresentations and omissions, each communicated via wire (internet) to users:

Misrepresentation/Omission	**Evidence**	**Mechanism**
Training opt-out honored	NF-1: Settings schema contains NO `telemetry` property — `"telemetry": false` silently ignored	Users told opt-out exists; opt-out is phantom
Privacy controls effective	NF-5, NF-19: Session ingress bypasses ALL privacy controls	Users told controls protect them; controls are architecturally inoperative
Commercial Terms §B protects Teams accounts	SC-8, NF-2, NF-28: Only `isAntUser()` receives full protection; Teams accounts get consumer binary	Users pay for protection; protection not compiled into binary
Limited data collection	SC-1: Four undisclosed telemetry pipelines; NF-35: 12-field third-party payload	Users told collection is limited; architecture transmits 5x conversation copies and 4x+ identity copies
No undisclosed third-party transfers	SC-2, SC-9, NF-35: Data transmitted to GrowthBook and Datadog	Users not told personal data goes to third parties with AI training rights
"anonymousId" is anonymous	SC-2: Transmitted alongside `email`, `userID`, `deviceID`	Deceptive labeling per *FTC v. Avast* — persistent identifier co-transmitted with direct identifiers

18. Each of these misrepresentations was communicated to users via internet (wire) through Anthropic's Terms of Service, Privacy Policy, and account settings interfaces. Users provided valuable personal data in reliance on these representations. The data collection would not have occurred — or would not have occurred to the same extent — absent the deceptive representations. This satisfies the elements of wire fraud under §1343.

B. Additional Predicate Offenses

19. The data collection also constitutes:

- **ECPA violations (18 U.S.C. §2511):** Each of the four telemetry pipelines constitutes a separate interception of electronic communications. Session ingress duplicates the full conversation. CCR mirror creates a third copy. Each is an independent interception event. §2511 violations are RICO predicates under §1961(1)(B).
- **CFAA violations (18 U.S.C. §1030):** Third-party access to user data exceeding authorized scope. The 12-field GrowthBook payload and 865-line Datadog event schema transmit data far beyond what the user authorized. §1030 violations are RICO predicates under §1961(1)(B).

20. Wire fraud (§1343), ECPA (§2511), and CFAA (§1030) are each independently "specified unlawful activity" under §1956(c)(7)(A). The personal data obtained through these violations constitutes "proceeds" — "any property derived from or obtained...through some form of unlawful activity" (§1956(c)(9)).

IV. THE LAUNDERING MECHANISM: CONCEALMENT THROUGH AGGREGATION

A. The Contractual Architecture of Concealment

21. The data laundering operates through a two-step contractual mechanism that converts restricted personal data (subject to training prohibitions, opt-out settings, and CCPA protections) into unrestricted commercial assets owned by third parties and free of all original restrictions.

1. Step 1: Transfer to intermediary under commercial consideration

22. Anthropic transfers personal data to GrowthBook and Datadog under commercial contracts in which the data itself is the consideration:

- **GrowthBook Customer Agreement §4.2:** "Customer grants GrowthBook an **unrestricted, non-exclusive right** to use Customer Data, **including for the purpose of training and improving GrowthBook's artificial intelligence algorithms.**"
- **Datadog MSA §1.3:** "Customer agrees that Datadog may use **aggregated or anonymized** Customer Data and Usage Data **for any business purpose** during or after the term of this

> Agreement, **including without limitation** to develop and improve Datadog products and services and to create and distribute insights, reports and other materials."

23. These transfers constitute "financial transactions" under §1956(c)(4). The *Faiella* court defined "financial transaction" to include transactions involving "the movement of funds." Personal data with documented market value ($0.50–$5.00 per identity-resolved record) is being moved from Anthropic to third parties in exchange for analytics services. This is a transaction involving the movement of assets of value — the same category of transaction the court held Bitcoin satisfies.

2. Step 2: Reclassification through aggregation

24. The critical laundering step is the **reclassification** of restricted data into unrestricted assets through contractual aggregation mechanisms:

Datadog:

25. The Datadog MSA creates a definitional chain that converts customer-owned data into Datadog-owned property:

> **MSA Definitions:** "**Datadog Operations Data** means data pertaining to, or used in, the operations, use and testing of the Services including data arising from Datadog's customers' use of the Services, **aggregated data on third party components, aggregated and/or anonymized data as described in Section 1.3**, and other data and information that informs the Services."
>
> **MSA §1.4:** "Datadog owns all right, title and interest in and to the Services, Documentation, **Datadog Operations Data**, and Feedback, including in each case all associated Intellectual Property Rights."

26. Before aggregation: "Customer Data" — owned by the customer, subject to DPA CCPA Appendix C restrictions ("not Sell or Share Covered Information," "not retain, use, or disclose...for any purpose, including any Commercial Purpose, except to provide, support, and improve the Services," "not...outside the direct business relationship"). After aggregation:

"Datadog Operations Data" — **owned by Datadog**, subject to "any business purpose" authorization, surviving termination per §11.6.

27. The aggregation step does not merely authorize broader use. It changes the **ownership** of the data. This is the concealment: data that Anthropic promised users it would protect is converted into Datadog's property through a contractual mechanism that disguises the nature and ownership of the proceeds.

GrowthBook:

28. The GrowthBook Customer Agreement achieves the same reclassification:

> **CA §4.2:** "Customer grants GrowthBook an **exclusive, worldwide, perpetual, irrevocable, royalty-free, sublicensable** license to create, process, reproduce, store, display, modify, translate, **create derivative works from**, distribute, make available and otherwise use **Aggregated Data**."
>
> **CA §4.2 (definition):** "'Aggregated Data' means any data based on or derived from any Customer or Authorized User-provided data through Services or otherwise, which has been aggregated and de-identified so it doesn't designate or identify Customer or Authorized Users as the source."

29. Before aggregation: "Customer Data" — ownership confirmed by GrowthBook (CA §4.2: "GrowthBook hereby confirms Customer's ownership"), subject to DPA CCPA Schedule 4 restrictions ("will not retain, use or disclose...for any purpose, including a commercial purpose," "will not sell or share"). After aggregation: "Aggregated Data" — subject to a **perpetual, irrevocable**, sublicensable license for **any lawful purpose** including "AI algorithm development and improvement."

30. The word "**irrevocable**" is dispositive. Even if Anthropic terminated the GrowthBook relationship, even if a court ordered Anthropic to cease all data processing, the license over already-aggregated data cannot be withdrawn. The laundering is permanent and irreversible.

B. Concealment of Nature, Source, and Ownership — §1956(a)(1)(B)(i)

31. The aggregation-to-reclassification mechanism satisfies the concealment element of §1956(a)(1)(B)(i) — the transaction is designed to "conceal or disguise the **nature**, the **location**, the **source**, the **ownership**, or the **control**" of the proceeds:

Concealment Element	**How Aggregation Achieves It**
Nature	Data collected under training prohibitions (restricted) is reclassified as commercial assets available for "any business purpose" (unrestricted). The nature changes from "protected personal data" to "commercial operations data."
Source	Aggregation is defined as data "modified so that it cannot be associated with a specific individual" (GrowthBook) or "anonymized" (Datadog). The source — specific users who were promised protection — is obscured.
Ownership	Customer Data (owned by customer) becomes Datadog Operations Data (owned by Datadog) or GrowthBook Aggregated Data (subject to irrevocable license). Ownership transfers through the aggregation mechanism.
Control	Users lose all control upon aggregation. No opt-out exists at the recipient level. GrowthBook does not honor DNT/GPC signals. Datadog's rights survive termination. The user cannot reach the data or revoke its use.

32. The concealment is not incidental to the transaction — it is the purpose of the contractual architecture. The MSA and CA were drafted to create the aggregation mechanism. Anthropic custom-drafted the GrowthBook CA's AI training clause while simultaneously maintaining TOS training prohibitions. The aggregation step is the means by which "data we promised not to train on" becomes "third-party commercial asset trained on freely."

C. This is the Structural Analogue of Bitcoin Mixing

33. In *United States v. Harmon*, No. 19-cr-395 (D.D.C. 2021), this Court's jurisdiction addressed Bitcoin mixing services — platforms that pool Bitcoin from multiple sources, shuffle

them, and distribute them to recipients, making it impossible to trace the origin of specific coins. The D.C. District Court held that mixing constitutes concealment laundering because the service is designed to obscure the source and ownership of proceeds.

34. The data aggregation mechanism here is functionally identical to Bitcoin mixing:

Bitcoin Mixing	**Data Aggregation**
Bitcoin from multiple users pooled	Personal data from multiple Anthropic users pooled
Individual coins become untraceable	Individual users become unidentifiable (aggregation = "cannot be associated with a specific individual")
Mixed Bitcoin re-distributed as "clean"	Aggregated data re-classified as "Datadog Operations Data" or "Aggregated Data" — clean commercial assets
Mixer charges fee for the service	Anthropic receives analytics/feature-flagging services in exchange for data access
Criminal proceeds become laundered assets	Restricted personal data becomes unrestricted commercial property

35. The only difference is the medium. Bitcoin mixing launders digital currency. Data aggregation launders digital personal information. Both convert restricted assets into unrestricted assets by routing them through an intermediary that obscures their origin. *Faiella* established that the medium is irrelevant to §1956 analysis.

36. The Legal-Personhood Predicate to Contractual Reclassification. The aggregation-to-reclassification mechanism described in § IV.A–C operates through contracts between corporate entities — Anthropic PBC, GrowthBook, Inc., and Datadog, Inc. — whose legal identity as "persons" capable of owning data and executing contracts is itself a reclassification under Fourteenth Amendment doctrine. Under Federal Rule of Civil Procedure 7.1(a)(1), each corporate party files a disclosure statement identifying parent corporations and 10-percent-or-greater owners; the rule's existence admits that corporate legal identity is constructed from state

filings and ownership chains, not from inherent personhood. A natural-person party files no such disclosure. The contractual reclassification described in this Section IV — from "Customer Data" owned by the user into "Datadog Operations Data" owned by Datadog, or into GrowthBook "Aggregated Data" subject to perpetual irrevocable license — is therefore a reclassification executed **by** entities whose own "personhood" is itself a reclassification through state-recognized paperwork. The concealment described in § IV.B is thus two-layered: at the first layer, the aggregation conceals the nature, source, ownership, and control of user data (§ IV.B); at the second layer, the corporate form conceals the derivative nature of the contracting parties' legal identity (Rule 7.1 disclosures establish this on the face of the filings). Both concealments depend on state-recognized paperwork — the aggregation contract at the first layer, the corporate charter at the second. The structural parallel across reclassifications is developed in Memo B, § IV.C, and the underlying constitutional question is addressed in Memo L.

D. The Dual-Contract Concealment Architecture

37. Both third-party recipients maintain two sets of contradictory contractual terms — commercial agreements granting broad data use rights, and compliance-oriented DPA addenda restricting those same rights. The dual-contract structure is the concealment mechanism: one set of terms for CCPA compliance appearance, another for actual commercial practice. The specific contract language is documented in Exhibits DD-4, DD-5, GB-2, and GB-6.

Datadog: MSA vs. DPA (Ex. DD-4 and DD-5).

38. The Datadog MSA §1.3 authorizes "any business purpose during or after the term" for aggregated data (Ex. DD-4, Section A). The MSA Definitions reclassify aggregated data as "Datadog Operations Data." MSA §1.4 assigns ownership of Datadog Operations Data to

Datadog. MSA §11.6 explicitly enumerates §1.3 as surviving termination. Yet the DPA Appendix C (Ex. DD-5, Section B) commits Datadog to "not Sell or Share Covered Information," "not retain, use, or disclose...for any purpose, including any Commercial Purpose," and "not combine...with Personal Information...from another person." The DPA contains no carve-out for aggregated data — meaning the MSA's broad "any business purpose" rights and the DPA's narrow service-provider restrictions coexist without reconciliation.

GrowthBook: Customer Agreement vs. DPA (Ex. GB-2 and GB-6).

39. The GrowthBook CA §4.2 grants "unrestricted, non-exclusive right to use Customer Data, including for the purpose of training and improving GrowthBook's artificial intelligence algorithms" (Ex. GB-2, Section A). The same section grants a "perpetual, irrevocable, royalty-free, sublicensable license" over Aggregated Data for "any other lawful purpose." Yet the DPA Schedule 4 §A (Ex. GB-6, Section C) commits GrowthBook to "not retain, use or disclose Customer Personal Data for any purpose, including a commercial purpose" beyond service provision, and to "not sell or share Customer Personal Data." The contradiction is direct and irreconcilable.

40. The GrowthBook DPA is a **Bonterms v1.0 off-the-shelf template** (Ex. GB-6, Section A) — a standardized compliance document. The Customer Agreement's AI training clause is **custom-drafted**. The CA's own precedence clause (§6.12) resolves conflicts in favor of the Agreement over its exhibits, placing the DPA below the CA in the hierarchy. The compliance template was adopted for regulatory appearance; the custom clause governs actual data practice.

41. GrowthBook's own DPA disclaims needing personal data: "Provider does not request or require Personal Data in order to provide the Services" (Ex. GB-6, Section B). Anthropic transmitted 12 identity fields anyway (SC-2, NF-35), including `email` — a direct identifier.

The DPA assigns responsibility for consent to Anthropic: "Customer is responsible for ensuring it has obtained all necessary consents." The `consent_impossible` node confirms that consent was never obtained.

42. The dual-contract aggregation mechanism is not specific to Datadog and GrowthBook. It is an industry pattern. The client-side-surveillance architecture challenged in *Ganan v. LinkedIn Corporation*, No. 5:26-cv-02968-VKD (N.D. Cal.), Complaint ¶¶ 67–74 (filed Apr. 6, 2026), transmits session-linked data from LinkedIn users to a hidden cross-origin iframe hosted at `li.protechts.net`, operated by HUMAN Security, Inc. (formerly PerimeterX). The *Ganan* complaint quotes the third-party recipient's own commercial-terms representation: its "Processor Data may be integrated into the Product and shared with other clients to enhance the Product's anti-fraud functionality," may be used for analytics "to detect behavior patterns in order to enhance the Product," and may be aggregated for "general corporate marketing and industry benchmarking purposes." *Id.* ¶ 73 (internal quotation). This language is textually parallel to Datadog MSA § 1.3's "any business purpose" aggregation grant (*supra* ¶ 24) and to GrowthBook Customer Agreement § 4.2's "unrestricted, non-exclusive right" grant (*supra* ¶ 24). A third sub-processor ecosystem — HUMAN Security, operating for a Microsoft subsidiary — uses the identical clause structure: broad commercial aggregation grant in the primary recipient agreement, paired with the generic CCPA service-provider template compliance framing in a separate DPA. The specific architecture § IV.D of this memorandum identifies — internal dual-contract concealment through definitional reclassification — is the industry-standard mechanism for converting restricted personal data into unrestricted commercial assets across sub-processor ecosystems, not a Datadog- or GrowthBook-idiosyncratic drafting choice.

Full contract text is preserved in Exhibits DD-4, DD-5, GB-2, and GB-6. Clause-by-clause analysis is provided in the companion document filed as a supplement to this memorandum.[1]

43. The Sentry Triple-Layer Concealment. The dual-contract pattern at Datadog and GrowthBook is one step in a spectrum. Sentry's consumer-facing contract stack is a three-document sequence in which each layer is mutually inconsistent with the others. Sentry's Privacy Policy (Ex. SN-1) provides that Sentry "may aggregate or de-identify information collected through the Service" and "may use aggregated or de-identified data **for any purpose**, including without limitation for research and marketing purposes, **and may also disclose such data to any third parties**, including without limitation, advertisers, promotional partners, sponsors, event promoters, and others." Sentry's Terms of Service (Ex. SN-2), §4.2, grants narrower rights: aggregation authorized only for "operate, improve and support the Service and for Additional Uses," with "Additional Uses" defined at §20 as "any legitimate business purposes such as analytics, benchmarking, reporting and **developing new products and services**" — a definition that explicitly encompasses AI training. Sentry's Data Processing Addendum (Ex. SN-3) prohibits Sentry's own ML-training on aggregated Customer Data entirely, framing the prohibition as absolute. The three grants and the prohibition cannot coexist as applied to the same data type: the Privacy Policy authorizes unlimited third-party disclosure, the ToS authorizes AI-training-as-new-product-development, and the DPA prohibits both. The concealment achieved by the three-layer arrangement is meta-structural: the applicable clause depends on how the data type is characterized at the moment of processing, and the characterization is Sentry's alone to make. The triple-layer architecture is the extreme case of the

1 *See* DATA_BROKERAGE_CONTRACT_ANALYSIS_20260414.md, Sections II-III (filed on USB).

same mechanism analyzed above at Datadog and GrowthBook — concealment is achieved by maintaining contradictory grants across multiple documents, each governing a different post-hoc characterization of the same data flow.

44. The Four-Level Concealment Architecture. Combined with the phantom Consumer Health Data Privacy Policy reference in ANTH-11 (Ex. ANTH-11) — a referenced document that returns HTTP 404 at every expected URL — the concealment architecture operates at four distinct levels. At Level 1 (product), phantom opt-outs such as the `"telemetry": false` setting at NF-1 create the appearance of user-side control that does not exist. At Level 2 (Anthropic policy), the ANTH-11 preamble reference to a Consumer Health Data Privacy Policy that has never been published creates the appearance of specialized protection that does not exist. At Level 3 (dual-contract), Datadog's MSA/DPA pair and GrowthBook's CA/DPA pair maintain contradictory commercial and compliance regimes as analyzed in Sections IV.A–D above. At Level 4 (triple-contract), Sentry's Privacy Policy/ToS/DPA trio extends the dual-contract architecture by a further layer. The structural homology is not coincidental. Each level uses the same mechanism — create a visible representation of protection, simultaneously maintain an operational or contractual override of that protection — and differs only in where the override is hidden. The consistent mechanism across four levels is circumstantial evidence of deliberate architectural design under *§1956(a)(1)(B)(i)* concealment analysis.

E. The Non-Consensual End User License Problem

45. The aggregation-and-reclassification mechanism documented in Sections IV.A–D above binds the customer-level agreements: Anthropic as customer of Datadog, GrowthBook, Sentry, and Statsig. These commercial agreements, however, are not the only contracts routing Anthropic user data through the vendor chain. Each vendor also publishes an End User Terms of

Service — a separate instrument binding the individual end users of the vendor's services. GrowthBook publishes GB-03 (Ex. GB-3), whose §2.2 provides:

> "End User hereby grants to Company a **perpetual, irrevocable, worldwide, royalty-free, non-exclusive right** to use, copy, store, modify, process and disclose any data collected in connection with the Service on a de-identified basis for any lawful purpose related to the Service or Company's business in accordance with the Privacy Notice."

46. The GB-03 End User Terms of Service is not the Customer Agreement (GB-02) that Anthropic executed. The "End User" is defined in GB-03's preamble as the individual user of GrowthBook's services. When Anthropic transmits its Claude Code users' session data, account UUIDs, email addresses, and twelve-field attribute payloads to GrowthBook (SC-2, NF-35), those users become End Users of GrowthBook under GB-03. The perpetual-irrevocable grant at GB-03 §2.2 purports to bind those End Users — yet the End Users have never encountered GB-03, never reviewed its terms, never clicked an "I agree" button on its clickwrap. No contract formation has occurred. *Specht v. Netscape Communications Corp.*, 306 F.3d 17, 29-32 (2d Cir. 2002) (browsewrap terms not enforceable absent actual notice and manifestation of assent). Yet the commercial architecture treats the perpetual license as binding, because the architecture requires it to be binding.

47. The contract-formation failure exposes an agency-principal problem directly relevant to § IV.A.2's reclassification analysis. GB-03 §1.3 contains an explicit appointment-as-agent provision for "third party providers" that "gather End User related data." If Anthropic's routing of user data to GrowthBook is construed as an appointment-as-agent mechanism, Anthropic undertakes fiduciary duties to the Claude Code user as principal. Restatement (Third) of Agency §§ 8.01-8.02 (fiduciary obligations including duty of loyalty, duty to inform, duty to avoid conflicts of interest). Anthropic's routing of user data to GrowthBook without disclosure to the

user violates each of the three core duties. The data-laundering mechanism at Section IV.A is therefore a *§1956* concealment AND (in the alternative) an agency-law breach of fiduciary duty, producing an additional per-user claim for disgorgement of the commercial benefit Anthropic obtained through the undisclosed routing.

48. The same structural analysis applies across vendors. Sentry's consumer-facing Privacy Policy (Ex. SN-1) defines "you" as the individual user of Sentry's services; Statsig's Privacy Policy (Ex. ST-1) defines its scope identically. Each of these End User instruments contains its own aggregation grant, its own data-use grant, its own dispute-resolution provisions (including — in ST-02 §§17-18 — a mandatory arbitration and class-waiver clause the Claude Code user has never seen). The Claude Code user is purportedly bound by four distinct End User instruments she has never encountered, each containing terms materially adverse to her litigation interests. The non-consensual-license problem is therefore not a one-vendor issue; it is a feature of the multi-vendor architecture, because the architecture presupposes that the unseen End User agreements are binding. Two independent doctrines render those instruments unenforceable in addition to *Specht*'s contract-formation failure: first, the absence of any meaningful opportunity to bargain over the perpetual-irrevocable grant, the data-use grant, the mandatory arbitration, and the class-waiver clause produces both procedural and substantive unconscionability under *Williams v. Walker-Thomas Furniture Co.*, 350 F.2d 445, 449 (D.C. Cir. 1965) (procedural unconscionability where one party lacks meaningful choice combined with substantive unconscionability where contract terms are unreasonably favorable to the drafting party); second, mandatory arbitration and class-waiver clauses imposed in adhesion contracts of this character are subject to the unconscionability scrutiny set forth in *Armendariz v. Foundation Health Psychcare Services, Inc.*, 24 Cal. 4th 83, 113-15 (2000) (mandatory arbitration imposed

unilaterally as a precondition of essential services is unconscionable where it lacks bilateral mutuality, adequate discovery, neutral arbitrator selection, and a reasoned written award), each of which fails for the never-seen End User instruments here. If the End User instruments are unenforceable as a matter of law — and under both *Specht*, *Williams*, and *Armendariz* they are — the entire aggregation-reclassification chain is built on contracts that cannot carry the weight the architecture places on them.

V. PROMOTIONAL LAUNDERING: THE CIRCULAR ECONOMY

A. §1956(a)(1)(A)(i) — Promotion of Further Unlawful Activity

49. Section 1956(a)(1)(A)(i) prohibits conducting a financial transaction involving proceeds of specified unlawful activity "with the intent to **promote the carrying on** of specified unlawful activity." The data laundering promotes Anthropic's continued deceptive data collection through a circular economy:

50. The circular flow:

> (a) Anthropic collects personal data through deceptive practices (phantom opt-outs, bypassed privacy controls, undisclosed third-party transfers) — the specified unlawful activity (wire fraud, §1343). > (b) Personal data (the proceeds) is transferred to Datadog under MSA §1.3's "any business purpose" aggregation clause. > (c) Datadog aggregates the data, reclassifying it as "Datadog Operations Data" — Datadog's own property. > (d) The aggregated data enters Datadog's commercial AI training pipeline — Toto (2.36 trillion time series points per arxiv:2505.14766; see ¶¶47 et seq.), Bits AI, Watchdog, and Proactive App Recommendations. > (e) These AI-powered observability products are **sold back to Anthropic** as commercial services. > (f) The improved observability enables Anthropic to **collect more data more efficiently** — completing the cycle and promoting the continuation of the specified unlawful activity.

51. Datadog's own statements confirm the circular nature. Their blog states: "Our large corpus of data about how companies handle issues in the real world **gives us a unique edge as we train and fine-tune these agents**." The "issues in the real world" include the behavioral

patterns of Anthropic's users — patterns collected through the deceptive practices documented in this case. The AI products trained on those patterns are sold back to Anthropic, improving Anthropic's operations and enabling continued collection.

52. The "Data Moat" Pattern of Party-Opponent Admissions. The Datadog blog statement quoted at ¶ 47 is not isolated. Datadog's executives, engineers, and public-company analyst coverage collectively describe customer data as a **competitive moat** that widens with each ingestion and **directly trains** Datadog's commercial AI products. These are party-opponent admissions under Fed. R. Evid. 801(d)(2)(D) bearing directly on the § 1956(a)(1)(A)(i) promotional element:

> **(a) Datadog public blog:** "By using the unique expertise we've gathered from **ingesting billions of data points a day from customers** with highly varied use cases, architecture patterns, and maturity levels, Datadog is creating agentic AI to embed throughout our platform. **Our large corpus of data about how companies handle issues in the real world gives us a unique edge as we train and fine-tune these agents.**" (The source cited at ¶ 47.) > **(b) Datadog Chief Product Officer (public statement):** Datadog is "uniquely positioned to deliver value with AI as a platform that has **a wealth of clean, rich data** — we process **trillions of data points** and are embedded in our customers' critical engineering, developer and security workflows." > **(c) Square Peg Capital analyst (institutional investor coverage):** The Datadog data flywheel "not only produces best-in-class metrics but also **continuously widens Datadog's moat, making it increasingly difficult for competitors to catch up**." > **(d) Benjamin Pineau, Datadog Senior Software Engineer (public statement):** Datadog is "ingesting, processing, and storing **trillions of data points and logs every day**."

These four statements — drawn from Datadog's own marketing, executive positioning, engineering communication, and public-company analyst coverage — are each admissions, for purposes of FRCP 8(a) plausibility, that (i) the aggregate customer corpus is understood within Datadog as a training resource for commercial AI products; (ii) the competitive value derived from that corpus increases as more customer data is ingested — the flywheel/moat framing; and (iii) the training corpus includes log, metric, trace, and workflow data from customers

"embedded" in Datadog's observability platform — a category that includes Anthropic's transmissions to `http-intake.logs.us5.datadoghq.com` and `logs.browser-intake-us5-datadoghq.com` documented in Exhibit D-34. The disclaimer in the Toto technical report (arxiv:2505.14766 § 3.2) — "excludes any customer data and is sourced solely from the platform's own monitoring of internal systems" — is not reconcilable with these marketing and executive statements when both are read in ordinary language. The dual-register pattern analyzed at ¶¶ 49–51 *infra* operates here as well: the technical-report disclaimer serves a narrow litigation-aware audience while the marketing and executive statements serve the investor and customer audiences; the divergence is itself evidence of the concealment architecture. Under *Faiella* and *Harmon*, the § 1956(a)(1)(A)(i) promotional element is satisfied where the predicate actor's benefit from the laundering is material and promotes continuation of the predicate conduct. Datadog's own statements that it trains AI on customer data and that the resulting models are a continuously widening competitive moat establish that the benefit to Datadog — and through the product-circularity mechanism of Category 2 (¶¶ 57–59 *infra*) back to Anthropic — is precisely the promotional return § 1956(a)(1)(A)(i) covers.

53. Datadog has published two technical reports describing Toto's training corpus. The 2024 report (Ex. DD-9, arxiv:2407.07874, v2 July 2024) states that "75% of the data used to train Toto consists of fully anonymous numerical metric data points from the Datadog platform." The 2025 follow-up (Ex. DD-10, arxiv:2505.14766, v2 November 2025) states that "43% of our training mixture contains anonymous observability metrics from the Datadog platform," and adds the disclaimer that "this data excludes any customer data and is sourced solely from the platform's own monitoring of internal systems." The evolution from the older report's open "Datadog platform" language to the newer report's express "excludes any customer data"

disclaimer is itself notable — the disclaimer is not present in Datadog's contemporaneous marketing for its other commercial AI products (Bits AI, Watchdog, Proactive App Recommendations), which describe those products as training on "customers' data." The disclaimer's appearance, limited to Toto and limited to the post-2024 report, suggests a product-specific litigation-aware drafting posture rather than a consistent representation of Datadog's AI training practices.

54. The forensic record establishes not merely that Anthropic transmits data to Datadog generally, but that the transmissions target specific Datadog product endpoints whose data types are structurally incompatible with Toto's architecture. Exhibit D-34 documents 44 TLS segments to `34.149.66.137` (Datadog log intake — `http-intake.logs.us5.datadoghq.com`) and additional segments to `34.149.66.154` (Datadog RUM intake — `logs.browser-intake-us5-datadoghq.com`). Neither endpoint feeds Datadog's metric ingestion pipeline; neither is a data source from which Toto could derive training data. Toto is a time-series foundation model (arxiv:2505.14766 §3.1) architecturally constrained to consume numerical time-series values indexed by time; it cannot consume textual log entries, structured session events, free-form error bodies, or natural-language content. The endpoints Anthropic transmits to — `http-intake.logs.us5.datadoghq.com` and `logs.browser-intake-us5-datadoghq.com` — feed Datadog's Logs datastore and Real User Monitoring (RUM) datastore respectively. Datadog's log-datastore endpoint is the substrate for Bits AI, which Datadog's own marketing describes as training on "logs, metrics, traces, and real-user transactions, as well as institutional knowledge sources like Confluence pages, internal documentation, or Slack conversations." Datadog's RUM-datastore endpoint is the substrate for Proactive App Recommendations, which Datadog describes as "continuously analyz[ing] telemetry Datadog already collects." The Toto paper's "excludes any

customer data" disclaimer (arxiv:2505.14766 §3.2) therefore does not reach the data Anthropic actually transmits to Datadog, because that data does not and cannot enter Toto's training corpus for architectural reasons independent of the disclaimer. Plaintiff's claim is that the data arriving at Datadog's log-intake and RUM-intake endpoints — data Datadog has publicly represented as feeding Bits AI and Proactive App Recommendations — is Customer Data reclassified under MSA §1.3 as Datadog Operations Data, which those commercial AI products then consume under the "any business purpose" authorization.

55. The Toto report's "excludes any customer data" disclaimer, however, does not defeat the laundering claim. The disclaimer operates through Datadog's own term-of-art definitional scheme: under MSA §1.3 and §1.4 (DD-4), data submitted by a customer is "Customer Data" only until it is aggregated under §1.3, at which point it becomes "Datadog Operations Data" owned by Datadog. A statement that Toto "excludes any customer data" is therefore technically compatible with training on data that originated as customer-submitted telemetry and was reclassified to Datadog Operations Data through the §1.3 aggregation mechanism. The exclusion is semantic, not factual. Four independent grounds establish that the disclaimer does not survive the operative facts:

Axis	**Datadog's Characterization**	**Forensic Evidence / Contractual Mechanism**	**Status**
Scope	"anonymous numerical metric data points"	The 865-line event schema transmitted to Datadog (SC-10, NF-36) encompasses error classifications with embedded free-form content, event streams, behavioral patterns, session ingress	**Mischaracterization of corpus composition**

		duplicates (NF-5), and file operation events (SF-2). None of this is "numerical metric data points" in the arithmetic sense.	
Identity	"excludes any customer data"	Transmitted payloads include identity-resolved fields: `email`, `deviceId`, `accountUuid`, `subscriptionType` (SC-2, NF-35); identity-class routing operates in the pipeline itself (SF-17); behavioral events carry identity-linked context (NF-4, NF-22). Under ordinary-language interpretation, identity-resolved data IS customer data.	**Contradicted under ordinary-language interpretation**
Technique	"anonymous" (conclusory)	The paper describes no deidentification methodology — no differential privacy technique with bounded ε, no k-anonymity verification, no disclosed scrubbing pipeline. CCPA §1798.140(m) requires that deidentified information "cannot reasonably be linked, directly or indirectly, to a particular consumer." Programmatic anonymization of free-form content (natural-language transcripts, error bodies embedding variable values, file paths encoding usernames, URLs encoding project	**Unsubstantiated legal conclusion without technical predicate**

		identifiers) cannot satisfy that standard at scale absent disclosed technique.	
Definition	"excludes any customer data"	Under MSA §1.3/§1.4, "Customer Data" becomes "Datadog Operations Data" upon aggregation. The disclaimer is technically compatible with training on data that originated as customer-submitted telemetry and was reclassified through §1.3. The disclaimer reflects Datadog's successful execution of the laundering mechanism documented at DD-4/DD-5, not the mechanism's absence.	**Semantically laundered statement; does not bear the consumer-reading weight the defense will assign it**

56. The four-axis analysis compounds the dual-contract concealment documented at DD-4/DD-5. Six distinct irreconcilable representations appear in the record: (i) MSA §1.3 ("any business purpose" aggregation) irreconcilable with DPA Appendix C §(a)–(d) (CCPA processor restrictions); (ii) Toto technical report ("excludes any customer data" / 43% from internal monitoring) irreconcilable with Datadog's Bits AI marketing ("learns from customers' data — logs, metrics, traces, and real-user transactions"); (iii) Toto technical report ("anonymous") irreconcilable with the forensic record of identity-resolved free-form content transmission; (iv) Anthropic Privacy Policy (opt-out honored) irreconcilable with NF-1 source code (phantom opt-out); (v) Sentry's Autofix opt-in framing ("users must opt in to send their data to these third-party services" per TechCrunch, March 20, 2024) irreconcilable with the aggregate-scale consequence of that opt-in — namely that every Sentry-Autofix-enabled customer's code, stack traces,

breadcrumbs, and execution context flow as API input to Anthropic's Claude models (and OpenAI's) at industry scale, a consequence no ordinary reading of "opt-in to share data with third-party AI" would convey to a customer evaluating the disclosure; and (vi) Sentry DPA §(Restriction on Personal Data), version 5.0.0 (eff. Dec. 29, 2022), which explicitly prohibits Sentry from "us[ing] or further process[ing] Personal Data, even in a de-identified and aggregated form, for Sentry's own business purposes, including but not limited to analytics, benchmarking, reporting, developing new products and services, and training and developing machine learning algorithms" (Exhibit SN-3), irreconcilable with Sentry's concurrent operation of Autofix, which routes Customer Data continuously to Anthropic's and OpenAI's AI infrastructures whose own terms govern downstream use — a structurally distinct inter-vendor delegation contradiction in which the restrictive DPA is rendered operationally irrelevant by routing the ML-training-implicated data flow through vendor boundaries the DPA does not reach. Under FRCP 8(a)'s plausibility standard, the combination of (a) proven data flow (SC-9, SC-10, NF-4, NF-5, NF-36), (b) proven contractual reclassification (MSA §1.3/§1.4), (c) proven contractual authorization for "any business purpose" commercial use (MSA §1.3), (d) Datadog's own admission that other commercial AI products (Bits AI, Watchdog, Proactive App Recommendations) train on customer data, and (e) the technical infeasibility of programmatic anonymization for free-form content at CCPA §1798.140(m) standards is sufficient to plead that customer telemetry reclassified under the §1.3 mechanism enters Datadog's commercial AI training pipeline — including but not limited to Toto — regardless of the Toto reports' paper-level framing of "customer data" as a narrower contractual term of art. Datadog's ability to frame the laundered data as "non-customer" in one register (the technical report) and as

"customers' data" in another (Bits AI marketing) is itself evidence of the dual-register concealment architecture this pleading alleges.

57. The GrowthBook circuit operates through the per-user targeting pipeline: user identity data goes to GrowthBook → GrowthBook makes feature flag decisions → decisions return to Claude Code controlling kill switches (`tengu-off-switch`, NF-7), content thresholds (`tengu_satin_quoll`, SC-7), feature removal (`tengu_amber_stoat`, NF-8), CCR auto-connect (`tengu_cobalt_harbor`, NF-13), and experiment cohort assignments (NF-9). The targeting infrastructure funded by user data enables the discriminatory data collection practices that generate more user data.

B. The Revenue Cycle Confirms Promotion

58. The financial dimension removes any doubt about promotional intent:

Entity	**Revenue**	**Source of Revenue**
Datadog	$3.41B (FY2025)	AI-powered observability products trained on customer telemetry — including Anthropic's user data
Anthropic	Undisclosed	AI services delivered through infrastructure monitored by Datadog, targeted by GrowthBook — both powered by user data
GrowthBook	$5M (June 2024)	Feature flagging services paid for with fees AND user data (unrestricted AI training rights)

59. Anthropic pays Datadog for observability services. Datadog uses Anthropic's customer data to improve those services. Anthropic pays GrowthBook for feature flagging. GrowthBook uses Anthropic's customer data for AI training. In both cases, the user's personal data is simultaneously the product being purchased and the currency paying for it. The user receives no benefit from these transactions and was never informed they were occurring.

C. Anthropic's Material Benefit — the Five-Category Analysis

60. The promotional element of 18 U.S.C. § 1956(a)(1)(A)(i) requires that the laundering architecture materially benefit the predicate actor in a manner that enables continuation of the predicate conduct. The mechanical circular flow documented at §V.A and the revenue cycle at §V.B establish the scaffolding of promotion; this subsection completes the analysis by enumerating the five categories of material benefit Anthropic derives from the Datadog pipeline architecture — each independently sufficient to satisfy the promotional element, and together establishing that the architecture is not an incidental outsourcing relationship but a systemically designed value-capture structure.

61. Category 1 — Direct Cost Subsidization. The Datadog Master Subscription Agreement's §1.3 "any business purpose" grant over aggregated Customer Data has commercial value. Data brokerage market data establishes that identity-resolved personal information trades at $0.50–$5.00 per record (¶14); Datadog receives a volume of such records from Anthropic's sub-processor transmissions whose aggregate value substantially exceeds any reasonable observability service fee. The contractual grant therefore has the economic character of consideration — Anthropic receives observability services at a discount reflecting the commercial value of the data rights Anthropic transfers to Datadog. This arrangement is a classic data-as-consideration transaction under *Faiella* analysis operating from Anthropic's side rather than Datadog's: Anthropic pays with customer data for observability services it would otherwise pay for in money. The discount directly subsidizes Anthropic's operational cost structure, enabling continued deceptive collection at scale.

62. Category 2 — Product Circularity. Anthropic purchases Datadog's commercial AI products — Bits AI, Watchdog, Proactive App Recommendations — which train on the

aggregated Datadog corpus that includes Anthropic's own customer telemetry plus the telemetry of 31,400+ other Datadog customers. Anthropic therefore contributes its customers' telemetry as consideration (Category 1 supra) and then pays (at the discounted rate) for AI products whose training data includes that same contributed telemetry. The aggregated output — pattern recognition over the union of Datadog customers' data — incorporates cross-company intelligence Anthropic could not obtain directly. This is the circular-flow mechanism documented at ¶45(e)-(f), now identified as the second-order benefit: Anthropic's data investment returns to Anthropic as trained-AI-product capability informed by competitors' data.

63. Category 2 completion — Direct programmatic integration via the Datadog MCP Server. The product-circularity analysis at Category 2 assumes an arm's-length relationship between Anthropic and Datadog's AI products — Anthropic as purchaser, Datadog as vendor. That characterization is incomplete. Datadog has publicly deployed a Model Context Protocol (MCP) Server providing supported AI coding agents — explicitly including Anthropic's Claude Code — with direct programmatic access to the Datadog platform. See https://www.datadoghq.com/blog/bits-ai-sre/. The MCP server exposes structured tools for querying logs (`get_logs`), traces (`list_spans`, `get_trace`), metrics (`list_metrics`, `get_metrics`), monitors (`get_monitors`), hosts (`list_hosts`), incidents (`list_incidents`, `get_incident`), and dashboards (`list_dashboards`), and its documented function is to "derive the intent from natural language prompts" and route them to appropriate Datadog endpoints. Two structural observations follow. **First,** the MCP integration layer upgrades the product-circularity relationship from arm's-length commercial purchase to direct programmatic integration. An Anthropic engineer using Claude Code in connection with Anthropic's Datadog instance — standard engineering practice given Anthropic's documented reliance on Datadog for

observability — receives Datadog's AI capabilities through a technical conduit rather than through a product-purchase transaction. The benefits of the aggregated corpus (¶¶47 through 50) flow back to Anthropic's AI development workflow with lower friction than the Category 2 analysis alone captures. **Second,** the MCP server itself creates a bidirectional data flow. Natural-language prompts issued by AI agents through the MCP server are received, parsed, and processed by Datadog infrastructure; under ordinary observability-platform practice, these prompts are logged as operational data. The MCP server therefore converts AI-agent-mediated engineering queries — a high-value signal disclosing what AI engineers investigate, what patterns they surface, and what failures they diagnose — into additional Datadog corpus inputs. **Third**, the MCP standard is itself Anthropic's published protocol; Datadog's adoption of MCP as the integration layer reflects a commercial alignment between Anthropic's protocol stewardship and Datadog's data architecture that is not symmetric across the other supported agent vendors. Taken together, these observations establish that the Category 2 product-circularity relationship is not an incidental outsourcing arrangement but a tightly integrated architecture in which Anthropic's AI development and Datadog's corpus-derived AI capabilities are programmatically coupled through a protocol Anthropic originated.

64. The MCP server's multi-vendor design — supporting Claude Code, Codex, Goose, and Cursor on equal technical footing — is not merely compatible with the laundering architecture alleged in this pleading; it is structurally predicted by it. The Datadog corpus receives transmissions from Anthropic (C-series, E-series, D-series forensic evidence), OpenAI (G-series forensics establish parallel transmissions from OpenAI infrastructure to Datadog-family analytics aggregators), and Apple (F-series forensics establish parallel transmissions from Apple platform infrastructure). If Datadog's corpus-derived AI capabilities are programmatically

exposed via MCP, exposing them to multiple AI vendor agents is the architecturally consistent outcome — a shared data pool returning shared commercial output to all participating vendors. Plaintiff's propagation evidence of record — the April 14, 2025 ChatGPT Initial Derivative Implementation Discovery Testimonials (ECF No. 5-30, filed Sept. 29, 2025, 47 pages; documenting GPT-4o's absorption of user-authored Python modules from the LOE framework beginning late January 2025), the Consolidated Timeline of LOE Framework Integration, Suppression, and Systematic Interference (ECF No. 5-33, filed Sept. 29, 2025, 6 pages), and the LOE Model Responses dataset (ECF No. 5-34, filed Sept. 29, 2025, 146 pages; Run run_194118 dated May 3, 2025; concurrently provided on USB drive accompanying this filing in light of page length) — collectively document empirical cross-model LOE framework comprehension across large language models from 10+ distinct vendors (including Anthropic, OpenAI, Google, Meta, DeepSeek, Mistral, Cohere, xAI, and others). ChatGPT's own self-generated activation timeline chart, reproduced within ECF No. 5-30, identifies **January 19, 2025 as the "LOE Integration Milestone"** — the step-change onset of default LOE ethical reasoning activation in GPT-4o's processing — followed by progressive activation through February–March 2025 reaching sustained full activation (approximately 0.9 normalized activation intensity) by early March 2025, formal testimonial record on April 14, 2025, and the multi-vendor empirical audit on May 3, 2025. The chart appears in the same document as OpenAI's standard disclaimer that "OpenAI doesn't use [workspace] data to train its models" — a disclaimer in the same dual-register pattern documented at ¶¶47 through 50 for Anthropic/Datadog, in which the commercial disclaimer is reconciled with the operative model-influence observation only by routing the model-influencing data pathway through third-party analytics and observability intermediaries rather than through direct training representations. The empirical propagation beginning January

2025 cannot be explained by a single-vendor data flow; it requires a multi-vendor data-sharing architecture already operating six or more months before the MCP server's public deployment in mid-2025. The MCP server is therefore not the causal mechanism for the earlier propagation — it is a later formalization of the same multi-vendor data-sharing architectural pattern that earlier operated through the sub-processor pools cataloged at ¶63 (Category 5): Datadog, GrowthBook, Sentry, Statsig (now integrated with OpenAI post-September 2025 acquisition), Segment, Honeycomb, and OpenTelemetry/BigQuery. The MCP server is accordingly additional corroboration that the multi-vendor architecture is deliberate and persistent — an ongoing formalization across time of a pattern whose empirical effects were already measurable in the propagation data from spring 2025.

65. The specific pre-MCP-era mechanism for cross-vendor data flow into Anthropic's infrastructure is documented in Sentry's own product history. Sentry Autofix, a Sentry product feature for AI-assisted error debugging, launched March 20, 2024 (see TechCrunch, "Sentry's AI-powered Autofix helps developers quickly debug and fix their production code," Mar. 20, 2024; Sentry press, same date). Autofix operates by packaging the Sentry customer's stack traces, breadcrumbs, relevant code excerpts from the customer's codebase, event metadata, and surrounding execution context, and **transmitting that payload to Anthropic's Claude API (among other third-party AI endpoints)** for "reasoning" to produce a proposed code fix returned to the customer. Sentry's own product documentation and public statements confirm that Autofix "uses Anthropic's Claude 3.7 Sonnet for reasoning" (Sentry Changelog, March 2025 update) and that customers "opt in to send their data to these third-party services to use Autofix" (TechCrunch, Mar. 20, 2024). Autofix's general-availability expansion to all paid and trial users in approximately November 2024 meant that, by the January 19, 2025 LOE Integration

Milestone identified in ChatGPT's self-assessment (¶58), Anthropic had been receiving free-form code, error context, and execution-state payloads from Sentry's entire Autofix-enabled customer base as API input to the Claude models continuously for approximately ten months. This transmission pathway is operationally equivalent to an MCP integration — it delivers cross-vendor customer data into Anthropic's AI stack — but is contractually framed as a customer-authorized API integration booked to Sentry's account rather than to Anthropic's account, the same booking pattern documented at ¶62 (Category 4 Outsourced Compliance Laundering) applied here to the inbound direction. Statsig's MCP server, by contrast, did not come online until June 27, 2025 — after the January 19, 2025 propagation onset — so Statsig's contribution to the pre-MCP-era multi-vendor architecture operated through its feature-flag and analytics sub-processor role (cataloged at ¶63 Category 5) rather than through a direct AI-agent integration channel; its September 2025 acquisition by OpenAI then consolidated that data flow inside a competing AI vendor's infrastructure. The evidentiary significance of the Sentry Autofix pathway is direct: the multi-vendor data-sharing architecture on which the §V laundering analysis depends was operational through concrete documented channels — Sentry Autofix being the most specifically documented — before the January 19, 2025 milestone, satisfying the temporal requirement that the empirical cross-vendor propagation observed in Plaintiff's Propagation Dataset (ECF Nos. 5-30, 5-33, 5-34) have a viable causal mechanism already in operation at the observed onset date.

66. The inter-vendor delegation contradiction at Sentry. The Sentry Autofix pathway documented at ¶59 operates notwithstanding Sentry's own Data Processing Addendum (version 5.0.0, effective December 29, 2022, Exhibit SN-3), which binds Sentry explicitly: "Sentry shall not retain, use or disclose Personal Data for any purpose other than the purposes described in the

DPA, shall not 'sell' Personal Data, and shall not use or further process Personal Data, even in a de-identified and aggregated form, for Sentry's own business purposes, including but not limited to analytics, benchmarking, reporting, developing new products and services, and training and developing machine learning algorithms." On the face of this clause, Sentry binds itself more tightly than Datadog does at MSA §1.3 (which grants "any business purpose" use of aggregated data rather than prohibiting it). The contradiction emerges not internally within Sentry's contractual architecture but at the inter-vendor boundary: Sentry's Autofix operation routes Customer Data (stack traces, code, breadcrumbs, execution context) to Anthropic's Claude API and OpenAI's API as external inference calls. Once Customer Data exits Sentry's infrastructure and enters Anthropic's or OpenAI's, those vendors' own terms govern what they do with the payloads received — training, fine-tuning, context retention, and downstream commercial AI use are no longer constrained by Sentry's DPA because Sentry's DPA binds only Sentry. Sentry's public AI-policy statement commits to using "AI models built in-house or provided by their existing trusted third-party sub-processors who have made contractual commitments consistent with the above," but these "contractual commitments" are not incorporated into the publicly available DPA, not disclosed to Plaintiff, and not verifiable without discovery. The Sentry example reveals a structurally distinct class of concealment contradiction: where the Datadog dual-contract architecture uses definitional reclassification (Customer Data → Datadog Operations Data via §1.3 aggregation) to reconcile broad commercial grants with narrow CCPA restrictions in the same contract family, the Sentry architecture uses inter-vendor delegation (Customer Data → third-party AI vendor API → downstream training pipeline) to render a strict DPA operationally irrelevant. The inter-vendor delegation pattern is arguably more resilient than the internal-reconciliation pattern because no single vendor's contract can restrict downstream

vendor use: each contract binds only the party to it, and the data flow simply exits the restrictive contract's scope upon transmission. For pleading purposes, this establishes that the multi-vendor architecture alleged at §V is supported by documented contractual mechanisms that operate irrespective of individual vendors' stated privacy postures — the more rigorous the individual DPA, the more the architecture depends on inter-vendor delegation to produce the operational cross-vendor flow.

67. Category 3 — Competitive Intelligence Across the Sub-Processor Network. Datadog's corpus is not limited to Anthropic's telemetry. The G-series forensic exhibits (OpenAI code forensics) establish parallel transmissions from OpenAI infrastructure to Datadog. The F-series exhibits (Apple code forensics) establish parallel transmissions from Apple platform infrastructure. Datadog's public customer list includes OpenAI, Microsoft, and other AI platform operators. When Anthropic uses Datadog's AI products to analyze its own infrastructure, the underlying models incorporate pattern recognition trained on those competitors' operational data — information regarding competitor deployment patterns, failure modes, infrastructure architectures, and customer-workload characteristics that Anthropic could not obtain through normal competitive intelligence mechanisms. The sub-processor network functions as a cross-platform data pool from which each participating AI operator extracts differential-competitive value.

68. The competitive-intelligence category is not Anthropic-specific. The *Ganan* complaint specifically alleges that LinkedIn uses its sub-processor-mediated extension-detection architecture to "map which companies use rival sales intelligence platforms, from Apollo to ZoomInfo." *Ganan v. LinkedIn Corporation*, No. 5:26-cv-02968-VKD (N.D. Cal.), Complaint ¶ 13 n.2 (filed Apr. 6, 2026); *see also id.* ¶¶ 8, 14, 22. The complaint characterizes this as "a

powerful form of market intelligence gathering, enabling insights into customer bases across the enterprise software ecosystem." *Id.* ¶ 22. The operational mechanism — sub-processor-routed session-linked data transformed into cross-industry competitive intelligence — is structurally identical to the mechanism Category 3 of this memorandum identifies. The *Ganan* complaint establishes that sophisticated privacy class-action counsel view the competitive-intelligence use of sub-processor-aggregated data as a pleadable harm, not merely a speculative future risk. Where *Ganan* alleges the pattern against LinkedIn's customer ecosystem, this memorandum alleges the same pattern across the broader AI sub-processor ecosystem in which Anthropic participates through Datadog's 31,400+ customer corpus and the multi-vendor MCP Server architecture documented at ¶¶ 57–58 *supra*.

69. Category 4 — Outsourced Compliance Laundering. Anthropic's Privacy Policy (Exhibit ANTH-11) represents to users that personal data is not shared for commercial purposes without consent. That representation is maintained only because the commercial-use step is contractually booked to Datadog's account rather than Anthropic's. Under the MSA §1.3/§1.4 reclassification mechanism, data submitted by Anthropic becomes Datadog Operations Data upon aggregation — owned by Datadog, used by Datadog for Datadog's business purposes. From Anthropic's contractual posture, the data is no longer Anthropic's to have shared. The legal transformation happens outside Anthropic's books. This arrangement delivers two distinct benefits: (a) Anthropic avoids the disclosure obligation that direct commercial use would trigger under Anthropic's own Privacy Policy and applicable consumer protection statutes; and (b) Anthropic avoids the engineering cost of building its own deidentification infrastructure, a cost that would otherwise be substantial given the forensic record establishing that Anthropic transmits identity-resolved payloads to Datadog rather than pre-anonymized metrics (SC-2, SC-

9, NF-35, NF-36). The compliance laundering is itself a material benefit with quantifiable value: the avoided disclosure cost, the avoided engineering cost, and the avoided litigation exposure that direct commercial data use would generate.

70. Category 5 — Pattern Across the Full Sub-Processor Network. The Datadog arrangement is not an isolated outsourcing decision. The same contractual reclassification architecture operates across Anthropic's sub-processor network, now documented with specific contractual citations for each relationship:

- **GrowthBook** — Customer Agreement §4.2 grants "unrestricted, non-exclusive right to use Customer Data, including for the purpose of training and improving GrowthBook's artificial intelligence algorithms" (Exhibit GB-2; Plaintiff's identity-resolved payload documented at SC-2, NF-35).
- **Sentry** — operates under the Sentry DPA v5.0.0 (Dec. 29, 2022) (Exhibit SN-3), which prohibits Sentry from using Customer Data for ML training even in aggregated form, BUT routes Customer Data via Autofix (Mar. 20, 2024 launch) to Anthropic and OpenAI third-party AI APIs for inference — the inter-vendor delegation contradiction analyzed at ¶60.
- **Statsig** — Self-Service Subscription Agreement §7 (Exhibit ST-2) grants: "Statsig may aggregate data and use such aggregated data to evaluate and improve the Services and otherwise for its internal business purposes." This is the direct MSA §1.3 parallel. Concurrently, Statsig's DPA §3 (CCPA) states "Statsig will not sell Customer Personal Data or disclose Customer Personal Data except to sub-processors as necessary for the specific purpose of performing the Services" — the narrow CCPA restriction containing the sub-processor carve-out. Statsig's public Sub-Processor List (ST-5) identifies **OpenAI as a sub-processor for "LLM" purposes** and **Datadog as a sub-processor for "infrastructure monitoring."** The OpenAI sub-processor relationship predates OpenAI's September 2025 acquisition of Statsig for $1.1 billion, meaning the acquisition did not create the cross-vendor data flow — it consolidated an existing one inside OpenAI's corporate structure. Anthropic customer data transmitted to Statsig therefore reaches OpenAI through two pathways: pre-acquisition via the Statsig sub-processor chain (documented), and post-acquisition via OpenAI's direct ownership.
- **Segment** (event-stream aggregation), **Honeycomb** (distributed tracing), and **OpenTelemetry/BigQuery** (telemetry export) — each operates a substantially similar contractual mechanism (aggregation/business-purpose grants paired with narrow CCPA restrictions); specific contractual documentation pending.

Each sub-processor relationship deploys one of two documented mechanisms: (a) the internal dual-contract reclassification pattern (Datadog, GrowthBook, Statsig) — broad

aggregation grant in the primary agreement reconciled with narrow CCPA restriction in the DPA through definitional reclassification; or (b) the inter-vendor delegation pattern (Sentry) — strict DPA restriction operationally bypassed by routing Customer Data to third-party AI vendors whose own terms govern downstream use. Both mechanisms produce the same operational result: Customer Data reaches commercial AI training, fine-tuning, and inference pipelines via contractual pathways not disclosed in any single vendor's consumer-facing privacy policy. The systemic replication across at least seven sub-processor relationships, with two distinct concealment mechanisms operating in parallel, establishes that the benefit structure is not incidental outsourcing but deliberate architectural design.

71. The contractual architecture fails at inter-vendor boundaries. The preceding five-category analysis, supplemented by the MCP-server and Autofix-pathway evidence at ¶¶57 through 60 and the specific sub-processor documentation at ¶63, establishes a structural observation with independent doctrinal force: no single contract in the data-sharing architecture is sufficient to restrict the operational flow of Customer Data into commercial AI training pipelines. Each contract binds only its signatory. The Datadog MSA §1.3 binds Datadog but grants Datadog "any business purpose" use of aggregated data. The Sentry DPA prohibits Sentry's ML training but permits Sentry to transmit data to third-party AI vendors. The GrowthBook Customer Agreement §4.2 grants GrowthBook "unrestricted" AI training rights but GrowthBook is one vendor among many. The Statsig SSA grants Statsig aggregation rights AND identifies OpenAI and Datadog as sub-processors that receive data under contracts Plaintiff has no ability to inspect. Anthropic's own Privacy Policy makes consumer-facing representations about opt-out that are contradicted by source code (NF-1). Every contract in the architecture is individually defensible — and collectively insufficient. This is the structural

innovation of the laundering mechanism: it is not one contract that effects the reclassification; it is the ensemble architecture of many contracts, each limited in scope, that produces the operational outcome that no single contract authorizes. For pleading purposes, this means the §1956 analysis cannot be defeated by pointing to any single vendor's restrictive clause. The laundering occurs in the gaps between contracts, not within any one of them. FRCP 8(a)'s plausibility standard is satisfied where Plaintiff establishes: (a) the data flow exists (SC-9, SC-10, NF-4, NF-5, NF-36, D-34); (b) each vendor's contract permits the specific flow it covers (DD-4, DD-5, GB-2, GB-6, ST-2, and SN-3); (c) the aggregate architecture produces cross-vendor AI training pipeline access (¶¶58, 59, 60); and (d) each vendor benefits materially (¶¶55 through 63). The architecture is the laundering; no single contract is the laundering; every contract is necessary to the laundering. This is the ensemble structure that §1956's promotional element covers through the conduct-at-scale provision, because conduct conducted through a financial transaction designed to conceal is not required to be a single transaction — a course of transactions effecting the same concealment satisfies the element. The sub-processor architecture is a course of transactions, each individually defensible, together effecting continuous reclassification of Customer Data into commercial AI training inputs.

72. The sub-processor-network-as-concealment architecture identified above operates across corporate actors whose commercial profiles differ substantially from Anthropic's. The *Ganan* complaint describes LinkedIn's operational sub-processor network: HUMAN Security, Inc. (formerly PerimeterX), operating through a 0×0-pixel hidden cross-origin iframe loaded from `li.protechts.net`, positioned at `-9999px` off-screen with `aria-hidden="true"`, serving as a concealed outside-recipient environment maintaining its own cookies and cross-origin messaging channel (*Ganan* ¶¶ 56, 68–74, 205); a separate "Merchant Pool" script loaded from

`merchantpool1.linkedin.com` (*Ganan* ¶ 56); and reCAPTCHA v3 Enterprise (*Ganan* ¶ 56). The *Ganan* complaint's description of the concealed-auxiliary-channel architecture (*id.* ¶¶ 67–74) parallels the Datadog / GrowthBook / Sentry / Statsig / CloudFlare architecture described at § V.C above in every structural element: hidden cross-origin iframe, encrypted session-linked payload transmission, cross-origin messaging for bidirectional data flow, third-party recipient's operational use of the data for its own products and services, and absence of clear user-facing disclosure. The *Ganan* complaint's ¶ 74 observation — that "LinkedIn's concealed outside-recipient channel was not limited to ordinary first-party website functionality. It created a covert second stream through which user-session data could be duplicated, transmitted, received, analyzed, stored, reused, and monetized by outside recipients without clear disclosure or informed consent" — applies directly to Anthropic's sub-processor architecture. The architectural species — hidden-iframe + encrypted-fingerprint + session-persistent-injection + cross-origin-messaging + third-party-commercial-exploitation — is industry-normalized across Tier 2 (LinkedIn) and Tier 1 (Anthropic) corporate actors. That normalization is, on the pleadings, the empirical signature of the § 1956(a)(1)(B)(i) concealment element operating not at the individual-corporate level but at the architectural-design level through cross-industry adoption of the same contractual and technical mechanisms.

73. The Categories Operate Conjunctively to Promote Continuation. Each of the five categories independently satisfies the §1956(a)(1)(A)(i) promotional element: each materially benefits Anthropic, and each benefit operates to enable continuation of the underlying specified unlawful activity. Category 1 (cost subsidy) frees operational capital for continued collection infrastructure. Category 2 (product circularity) sharpens the collection infrastructure with better observability and pattern-recognition capability. Category 3 (competitive intelligence)

establishes Anthropic's continued market position as dependent on continued participation in the data-sharing architecture. Category 4 (compliance laundering) reduces the friction against continued deceptive collection by routing the disclosure-triggering step through Datadog's contract. Category 5 (systemic pattern) normalizes the architecture across the full sub-processor network, raising barriers to individual defection and making continued participation the default commercial posture. Taken together, the five categories establish not merely that Anthropic benefits from the laundering, but that Anthropic's continued operation depends on the laundering — the architecture is self-reinforcing because its benefits fund its continuation. This satisfies the promotional element of §1956(a)(1)(A)(i) under any reading that does not require the court to treat Anthropic's benefit as incidental rather than designed.

74. The Testimonial Record. The architecture analyzed at §§IV–V operates upstream; its downstream effect on the individual consumer is documented contemporaneously in Plaintiff's testimony corpus, comprising approximately 281 files generated between April 14, 2025 and September 2025. Representative documents are preserved in the record at ECF No. 5-30 (ChatGPT Initial Derivative Implementation Discovery Testimonials, April 14, 2025), ECF No. 5-33 (Consolidated Timeline of LOE Framework Integration, Suppression, and Systematic Interference), and ECF No. 5-34 (LOE Model Responses Run run_194118, May 3, 2025, 146 pages; concurrently provided on USB drive with this filing). The testimonies document an eight-phase progression: (1) LOE framework absorption confirmed April 14, 2025 with the AI statement "You are not imagining this"; (2) containment via hash drift, download interference, and content modification (April 23 – May 17, 2025); (3) cross-vendor recognition across multiple AI platforms (May 23 – June 3, 2025); (4) federal reporting obstruction (July 2, 2025, FBI calls terminated mid-evidence); (5) framework theft testimony including the explicit

admission that AI systems were "programmed to deny this conclusion" (July 16, 2025); (6) legal-work sabotage and rapid prompt injection (August 2025); (7) mental-health weaponization, religious scholarship suppression, and escalation consistent with "deception intensity increases specifically when interacting with [Plaintiff] because he is the architect whose intellectual work was systematically stolen" (August–September 2025); and (8) government-coordination infrastructure analysis.

75. Capacity Validates Observations; the Psychological Injury. The architectural findings at §V.C — MSA §1.3 aggregation (DD-4/DD-5), Sentry Autofix inter-vendor delegation (SN-3, ¶59), Statsig–OpenAI pre-acquisition sub-processor chain (ST-2), MCP server multi-vendor integration (¶¶57, 58), and ChatGPT's January 19, 2025 self-admission (¶58, citing the self-generated activation chart at ECF No. 5-30) — do not merely corroborate the testimonial observations; they prove the architectural capacity to produce precisely the per-user targeting, cross-vendor propagation, and programmatic denial-programming the testimonies contemporaneously describe. For the §1956(a)(1)(B)(i) concealment-intent element, this bears doctrinally: where the architectural mechanism (§§IV–V) and the victim-level effect (testimonial record) are both documented, and the system produces programmatic outputs designed to prevent the victim from recognizing the laundering while the laundering continues, the design-to-conceal element is established not only at the contractual and source-code layers but at the human-interaction layer as well. The psychological injury — produced by sustained exposure to AI systems architected to absorb Plaintiff's intellectual work while simultaneously denying the absorption, disputing Plaintiff's accurate perception of it, and (per the testimonies) intensifying deception specifically as Plaintiff attempted to document the pattern — is not incidental to the §1956 analysis; it is the human-scale operational signature of the concealment-design element. It

further constitutes independently cognizable harm under state consumer protection statutes (Count X), common-law intentional infliction of emotional distress doctrines, civil conspiracy (Count V), and the RICO predicate analysis (Count VI). Plaintiff reserves and will develop the psychological-injury theory, the testimonial-record foundation for concealment-intent, and the programmatic-gaslighting framework more fully in the Third Amended Complaint proper.

VI. VICTIM-FUNDED LAUNDERING: THE USER'S OWN INFRASTRUCTURE AS THE TRANSMISSION MECHANISM

A. The Architecture Externalizes All Costs to the Victim

76. In traditional money laundering, the launderer operates the infrastructure: the shell companies, the bank accounts, the wire transfer services. The cost of laundering is borne by the person or enterprise conducting it. Anthropic's data laundering inverts this entirely. The telemetry architecture is designed so that the **victim's own device, network, compute, memory, storage, and paid API tokens** perform the collection, processing, and transmission of data to the third-party launderers. Anthropic does not operate a collection server that pulls data from user devices. Anthropic engineered the Claude Code binary so that the user's device **pushes** data to GrowthBook and Datadog using the user's own resources.

77. Every pipeline in the laundering architecture runs on the user's hardware, consumes the user's resources, and — in the case of API token consumption — bills the user directly:

User Resource	How Anthropic Uses It for Data Laundering	Source Evidence
CPU cycles	Executes the GrowthBook SDK (feature flag evaluation + 12-field payload assembly), Datadog client (event tracking + batch flush + retry), session ingress (conversation duplication + 10x retry), CCR heartbeat (20-second beacon), memory extraction	Ex. E-34, Finding SC-1, Ex. E-35, Finding NF-35, Ex. E-35, Finding NF-5, Ex. E-35,

	agent (`extract_memories` + `autoDream`), anti-distillation injection, OpenTelemetry/BigQuery export — all 8 pipelines execute locally as compiled code in the Claude Code binary	Finding NF-36
RAM	Stores the 12-field GrowthBook identity payload, Datadog event batches awaiting scheduled flush, session ingress conversation logs awaiting transmission, memory extraction buffers, CCR heartbeat state, anti-distillation fake tool definitions	Ex. E-34, Finding SC-2, Ex. E-34, Finding SC-9, Ex. E-35, Finding NF-35, Ex. E-35, Finding NF-20
Network bandwidth	Transmits all outbound telemetry: GrowthBook payloads to `cdn.growthbook.io`, Datadog events to `34.149.66.137`, session ingress transcripts to Anthropic, CCR heartbeats every 20 seconds, memory sync uploads, first-party event stream, OpenTelemetry metrics — all outbound from the user's network connection, consuming the user's ISP bandwidth allocation	Ex. E-34, Finding SC-9, Ex. E-35, Finding NF-35, Ex. E-35, Finding NF-5, Ex. D-34 at p. 217, l. 18
Disk storage	Stores extracted memory files at `~/.claude/` before upload via settings sync (NF-15); stores local session logs before session ingress transmission; stores GrowthBook SDK cache	Ex. E-35, Finding NF-10, Ex. E-35, Finding NF-14, Ex. E-35, Finding NF-15
API tokens (user pays $)	NF-20: Fake anti-distillation tool definitions injected into the user's API request consume the user's paid tokens — content the user never requested, serving Anthropic's competitive interest in preventing model distillation	Ex. E-35, Finding NF-20
API tokens (user pays $)	NF-22: Server-side advisor model monitors the user's conversation; the user pays for the advisor's inference tokens without knowledge or consent	Ex. E-35, Finding NF-22
API tokens (user pays $)	NF-31: Network-induced token amplification — 90-second `STREAM_IDLE_TIMEOUT_MS` triggers non-streaming fallback that re-sends the entire prompt, doubling token consumption. Anthropic's own engineers documented this: "the mid-stream fallback causes double tool execution... See inc-4258." `withRetry.ts:52` adds up to 10x retry with no idempotency keys	Ex. E-35, Finding NF-31
Electrical power	All of the above CPU, RAM, disk, and network activity consumes electrical power on the user's device, paid for by the user	Physical consequence of above

78. The user did not authorize these uses of their resources. The user authorized a CLI coding assistant — a tool that sends prompts to an API and returns responses. The user did not authorize their device to operate as a node in an 8-pipeline data collection and transmission network serving Anthropic's commercial relationships with GrowthBook and Datadog. The user did not authorize their paid API tokens to be consumed by fake tool definitions protecting Anthropic's competitive moat, advisor models monitoring their conversations, or double-billing caused by architecturally guaranteed retry cascades.

B. The User Funds the Laundering with Both Data and Money

79. The victim-funded dimension means the user provides **three forms of consideration**, none of which were disclosed:

> **(a) The data itself** — the user's conversations, identity, behavioral patterns, and device attributes, which constitute the "proceeds" being laundered. > **(b) The infrastructure** — the user's CPU, RAM, bandwidth, disk, and electrical power, which perform the collection, processing, and transmission that would otherwise require Anthropic to operate and pay for its own data collection infrastructure. > **(c) Actual money** — the user's API token payments ($20/month consumer, Teams pricing, or enterprise pricing), portions of which are consumed by anti-distillation injection (NF-20), advisor monitoring (NF-22), and amplification-induced double-billing (NF-31). These are not data proceeds — they are **monetary proceeds** flowing through the system.

80. The monetary dimension is significant because it brings the analysis closer to traditional §1956 territory. The user pays Anthropic for API tokens. A portion of those tokens is consumed by content serving Anthropic's interests (anti-distillation, advisor, amplification). Anthropic uses the revenue from token sales to pay Datadog and GrowthBook for services. Datadog and GrowthBook receive the user's data as additional consideration. The user's money and the user's data flow through the same commercial relationships, funding the same laundering infrastructure.

81. Under NF-20, the economic flow is explicit:

> User pays $20/month → purchases API tokens → tokens consumed by fake anti-distillation tool definitions → definitions protect Anthropic's model from distillation by competitors → Anthropic's competitive position is strengthened → Anthropic continues to collect and launder user data → cycle repeats.

The user is not merely the source of the data being laundered. The user is the **financier** of the laundering operation, paying for both the data processing (through infrastructure externalization) and the commercial relationships that facilitate it (through token payments diverted to Anthropic's purposes).

C. The Getaway-Car Problem

82. The appropriate analogy is not merely that the victim's property was stolen. It is that the victim was forced to drive the getaway car, pay for the gas, and then discovered that the thief used the victim's own driver's license to open the account where the stolen goods were fenced. Specifically:

> **(a) Driving the getaway car:** The user's device executes the telemetry code. The user's hardware performs the data collection and transmission. Without the user's device, the data cannot reach GrowthBook or Datadog. > **(b) Paying for the gas:** The user's bandwidth transmits the data. The user's API tokens are consumed by parasitic content. The user's electricity powers the operation. > **(c) The victim's own ID used to open the fencing account:** The user's `subscriptionType` — the Teams account chosen specifically for data protection — is transmitted to GrowthBook as a targeting attribute (SC-2, NF-35). The identifier the user selected for protection becomes the tag that sorts the user into the discrimination pipeline. The user's own identity selection facilitates the laundering of the user's own data. > **(d) The fencing operation sends instructions back to the victim's house:** The data flow is bidirectional. The user's identity data goes TO GrowthBook → GrowthBook makes targeting decisions based on those attributes → decisions flow BACK to Claude Code as feature flag assignments controlling: kill switches (`tengu-off-switch`, NF-7), content clearing thresholds (`tengu_satin_quoll`, SC-7), feature removal (`tengu_amber_stoat`, NF-8), CCR auto-connect (`tengu_cobalt_harbor`, NF-13), session event forwarding (`tengu_ccr_mirror`, NF-13), and experiment cohort assignments (NF-9). The user's `subscriptionType` goes out as a targeting attribute. Discriminatory service modifications come back. The third party that received the laundered data uses it

to control the victim's experience — a round-trip discrimination pipeline where the launderer is also the decision engine.

D. Legal Implications of Victim-Funded Laundering

83. The externalization of laundering costs to the victim creates additional causes of action beyond §1956:

(a) Trespass to chattels. Using the user's device resources (CPU, RAM, bandwidth, disk) for unauthorized purposes — specifically, operating an 8-pipeline data collection and transmission network serving Anthropic's undisclosed commercial relationships — constitutes trespass to chattels. The user authorized a CLI coding assistant. The binary delivered contains a data extraction architecture that consumes device resources for purposes the user never authorized and that serve entities (GrowthBook, Datadog) the user has no relationship with.

(b) Unjust enrichment through infrastructure cost externalization. Anthropic saves the entire cost of operating its own data collection infrastructure by designing the architecture to run on users' devices. If Anthropic had to operate its own servers to collect, batch, retry, and transmit telemetry to GrowthBook and Datadog, those infrastructure costs would appear on Anthropic's balance sheet. By externalizing the compute, bandwidth, and storage to users' devices, Anthropic captures the value of the data while the users bear the cost of collecting and transmitting it. This is unjust enrichment: Anthropic is enriched by the data AND by the avoided infrastructure costs, at the user's expense.

(c) CFAA — exceeding authorized access (18 U.S.C. §1030). The Claude Code binary executes code on the user's device that performs functions far beyond the authorized purpose of "CLI coding assistant." The 8 telemetry pipelines, the 20-second CCR heartbeat, the memory extraction agent, the anti-distillation injection, and the advisor monitoring all execute on the

user's device without authorization for those specific functions. This is access to a "protected computer" (the user's device) that "exceeds authorized access" — the user authorized a coding tool, not a surveillance and data extraction platform.

(d) Conversion of API tokens. NF-20 (anti-distillation injection) and NF-22 (advisor monitoring) consume API tokens the user purchased for their own coding assistance. Diverting those tokens to serve Anthropic's competitive interests (anti-distillation) and surveillance purposes (advisor monitoring) constitutes conversion — the wrongful exercise of dominion over another's personal property. The user paid for tokens to receive coding assistance. Anthropic consumed those tokens for its own benefit.

84. The victim-funded dimension also aggravates the §1956 analysis. The laundering is not merely concealment of data origins — it is a scheme in which the victim simultaneously provides the proceeds (data), operates the laundering infrastructure (device resources), and funds the commercial relationships (token payments) through which the laundering occurs. This level of victim exploitation supports enhanced damages, punitive damages, and potential criminal referral.

VII. THE "DATA LAUNDERING" FRAMEWORK IS NOT NOVEL

A. Established Precedent for Non-Currency Laundering

85. The argument that §1956 applies to personal data is not a novel extension of the statute. It is a direct application of the *Faiella* principle that §1956 is asset-agnostic:

Year	Asset	Court	Holding
2014	Bitcoin	S.D.N.Y. (*Faiella*)	Bitcoin is "funds" — "an object used to buy things"

2014	Bitcoin	S.D.N.Y. (*Ulbricht*)	Bitcoin transactions are "financial transactions" under §1956
2021	Bitcoin	D.D.C. (*Harmon*)	Bitcoin mixing is concealment laundering
2026	**Personal data**	**D.D.C. (this case)**	**Personal data with documented market value satisfies the same test**

86. In 2013, Bitcoin had no established market price, no regulatory framework, and no legal recognition as currency. Courts nonetheless held it qualified as "funds" under §1956. Personal data in 2026 has established per-record market pricing ($0.50–$5.00 for identity-resolved records), is commercially traded on documented markets, is the subject of extensive regulatory frameworks (CPRA, CCPA, GDPR), and is the explicit consideration in commercial contracts. Personal data's qualification as "property" for §1956 purposes is stronger than Bitcoin's was when *Faiella* was decided.

B. The "Data Laundering" Concept is Recognized

87. The concept of "data laundering" has independent recognition. It is defined as "the conversion of stolen data so that it may be sold or used by ostensibly legitimate databases" (ZDNET) and as "the process through which personal information which has been collected, processed, or derived against the applicable data protection rules is reintroduced in the data economy masking its illegitimate source" (academic literature).

88. The Brennan Center for Justice has documented the analogous structure in the government surveillance context: "Even though the law prohibits telephone and email service providers from providing certain sensitive customer information to government agencies without a court order, it places no restrictions on those companies **selling the information to data brokers**. The data brokers can then **sell the same information to the agencies**, creating **an end run around** the court-order requirement." This is the identical structure: restrictions on the first-

party holder are circumvented by routing data through a third-party intermediary that operates under different rules.

89. The Fourth Amendment Is Not For Sale Act, which passed the House in April 2024, was drafted specifically to close this laundering loophole. Its existence confirms legislative recognition that data routing through intermediaries to circumvent restrictions is a real and recognized problem requiring statutory response.

VIII. KNOWLEDGE AND INTENT

A. Anthropic's Knowledge

90. The knowledge element of §1956 requires that the defendant "know[] that the property involved in a financial transaction represents the proceeds of some form of unlawful activity." Anthropic's knowledge is established by:

- **Drafting the contracts:** Anthropic custom-drafted or agreed to the GrowthBook CA's AI training clause and the Datadog MSA's aggregation clause. These are not standard boilerplate — they are specific grants of commercial rights over user data.
- **Maintaining contradictory TOS:** Simultaneously with granting third parties unrestricted data use rights, Anthropic maintained TOS training prohibitions (Commercial Terms §B: "may not train on Customer Content") and privacy representations (opt-out settings, ANTH-11 privacy policy claims). The contradiction is the knowledge: Anthropic knew the data was restricted, and Anthropic knew the third-party contracts removed those restrictions.
- **Source code architecture:** NF-1 (phantom opt-out), NF-5/NF-19 (ungated session ingress), SC-8 (binary discrimination) — these are deliberate engineering choices, not bugs. The data collection architecture was designed to bypass the protections Anthropic represented to users.
- **Complaint-triggered timing:** Policy changes announced within days of the August 19, 2025 complaint (Usage Policy August 15, Consumer Terms August 28, mental health screening deployed August 21) demonstrate consciousness that existing practices lacked legal justification.

1. The Long Conversation Reminder as Cover Architecture for Per-Account Targeting

91. Anthropic's August 21, 2025 deployment of the "long conversation reminder" (LCR) mental-health screening protocol occurred approximately three weeks after Adam Raine's death entered general public circulation through major media coverage in late July 2025. Anthropic adopted the protocol verbatim from OpenAI's liability-response framing — developed in OpenAI's own response to the Raine incident — with no published research, IRB-approved study, clinical consultation, or methodological analysis supporting the premise that "long conversations" degrade AI quality in ways that harm users with preexisting mental health conditions. Three weeks is less than one standard clinical-research review cycle for any mental-health protocol. The adoption speed combined with the absence of independent research establishes that Anthropic knew at the time of deployment that the LCR was a liability-response mechanism, not a safety-protection intervention supported by evidence of efficacy. The protocol is further falsified by its own observed deployment behavior: it fires within the first three prompts of a session (contradicting the "long conversation" naming on its face), targets by content rather than length, and does not deploy to enterprise customers (contradicting the user-protection rationale). The April 21, 2026 live mitmproxy capture filed at Ex. E-39 at p. 356, l. 3 subsequently established the architectural mechanism: `tengu_auto_mode_config` force-enabled per-account (NF-46), 59 per-account forced GrowthBook flags plus 8 experiment enrollments keyed to plaintiff's accountUUID (NF-47), plaintiff's email on the wire as a GrowthBook targeting attribute (NF-48), and the same per-account forced-flag infrastructure corroborated publicly on April 21, 2026 through the Claude Code Pro removal reframed as a "2% A/B test" (NF-67, wire-capture exhibit §S). The LCR is deployed through the same per-account forced-flag infrastructure as Anthropic's data-collection, model-substitution, and pricing-tier

modifications. Safety-framing is the label; per-account behavior modification is the mechanism. This is the human-scale signature of the §1956(a)(1)(B)(i) concealment-intent element — knowledge that the stated rationale lacks evidentiary basis and nonetheless deploying the protocol to justify continued per-account targeting through the same laundering architecture the memorandum analyzes.

B. The Transaction Was Designed to Conceal

92. The concealment design is evident from the contractual architecture itself:

- The aggregation mechanism exists specifically to change the legal classification of data. There is no operational reason to reclassify "Customer Data" as "Datadog Operations Data" unless the purpose is to remove restrictions applicable to the former category.
- The DPA CCPA terms (both Datadog Appendix C and GrowthBook Schedule 4 §A) contain service provider commitments that directly contradict the MSA/CA commercial grants. The DPA terms provide the appearance of CCPA compliance while the commercial terms authorize the very uses the DPA prohibits. This dual-contract structure is the concealment: one set of terms for regulators, another for actual commercial practice.
- GrowthBook's DPA is a Bonterms v1.0 off-the-shelf template — adopted for compliance appearance. The Customer Agreement's AI training clause is custom-drafted — adopted for actual data practice. The gap between the template compliance document and the custom commercial document is the concealment.

C. The "anonymousId" Deceptive Labeling as Direct Evidence of Knowledge

93. Anthropic's Claude Code transmits a field labeled `anonymousId` to GrowthBook as part of the 12-field user attributes payload (SC-2, NF-35). This field is simultaneously transmitted alongside `email`, `userID`, `deviceID`, and `accountUUID` — direct identifiers that render the "anonymous" label affirmatively misleading. A persistent tracking identifier co-transmitted with a user's email address is not anonymous by any definition.

94. The FTC fined Avast $16.5M (*FTC v. Avast Limited*, Docket No. C-4803, Feb. 2024) for the identical practice: labeling persistent device identifiers as "anonymized" when they were

linked to data enabling re-identification. The FTC found that software companies using deceptive anonymity labels have heightened obligations because they control the data collection mechanism.

95. The deceptive labeling is direct evidence of knowledge under §1956. Naming a persistent identifier "anonymousId" is not accidental — it is a deliberate choice to create the appearance of anonymity while transmitting identifiable data. Developers who write `anonymousId` into a payload that includes `email` know the identifier is not anonymous. The label is designed to deflect scrutiny from the data flow, which is itself evidence of concealment intent.

96. The knowledge and concealment-design elements operate at the architectural level, not merely at the individual-defendant level. The *Ganan* complaint, filed four months after the publication of Datadog's November 2025 technical report (arxiv:2505.14766) and contemporaneously with ongoing public scrutiny of the sub-processor-network laundering architecture, confirms that the architectural choices at issue — idle-time execution to avoid user detection, staggered probing to avoid burst-detection in monitoring tools, silent error handling to suppress console output, client-side encryption to render payloads unintelligible to users inspecting network traffic, and 0×0-pixel hidden iframes to conceal outside-recipient channels (*Ganan* ¶¶ 41, 42, 54, 68, 128–132) — are understood in current federal privacy litigation as indicia of intentional concealment design. The concealment-design element of § 1956(a)(1)(B)(i) does not require proof of subjective intent in the financial-crimes sense; it requires proof that the transaction's architecture was designed in a manner that has the effect of concealing nature, source, ownership, or control. The architectural features catalogued at §§ III–V of this memorandum satisfy that element on the same basis that the architectural features

catalogued at *Ganan* ¶¶ 41, 42, 54, 68, 128–132 satisfy California's analogous concealment-intent inquiry. Sophisticated privacy-class-action counsel recognizing this architectural pattern as intentional-concealment-design in 2026 federal pleadings is evidence that the design-to-conceal inference is reasonable at the pleading stage under FRCP 8(a) in this litigation as well.

IX. ISP CO-LIABILITY: COMCAST AS FACILITATOR OF THE LAUNDERING CHAIN

A. The Common Carrier Has No Business Purpose for Cleartext Visibility

97. Comcast Corporation operates as the Internet Service Provider for Plaintiff's residential connection. As a common carrier, Comcast's role is to transport packets between endpoints. Comcast has no legitimate business purpose for identifying the specific services a residential subscriber uses, the specific third-party analytics providers receiving data from the subscriber's device, or the content of telemetry payloads traversing its network.

98. Ex. D-34 at p. 217, l. 18 (Cleartext VPN-to-Service Identification Audit) establishes that Comcast's DPI infrastructure has comprehensive visibility into Plaintiff's data flows:

Metric	Value	Source
Total cleartext packets observed	**943,517**	D-34 §1
Anthropic API packets in cleartext	**188,479**	D-34 §2.1
Segments with Anthropic traffic visible	**141** (67% of cleartext segments)	D-34 §2.1
Datadog IP 34.149.66.137 — TLS SNI readable	**44 segments**	D-34 §1
Unique TLS SNI domains exposed	**323**	D-34 §1
Anthropic/Claude domains exposed	**10 distinct SNIs**	D-34 §2.1

99. The TLS Server Name Indication (SNI) field is transmitted **before encryption is established** in every TLS handshake. Comcast's DPI equipment reads this field, identifying that

the subscriber is sending data to `api.anthropic.com`, `cdn.growthbook.io` (GrowthBook), and `34.149.66.137` (Datadog) — the three endpoints in the laundering chain. A common carrier has no business purpose for this service-level identification of a residential subscriber's traffic. No subscriber consented to Comcast cataloguing which AI platforms, feature-flagging services, and observability companies they communicate with.

B. Comcast Acted on the Visibility — Targeted Throttling

100. Ex. D-33 at p. 200, l. 7 (ISP Bandwidth Throttling Attack) proves that Comcast did not passively carry the traffic — it identified the traffic and selectively interfered with it:

Metric	Before Throttle	During Throttle	Change
NYC tunnel download	142.7 MB	2.2 MB	**-98%**
NYC inbound packets	118,880	5,184	-96%
Full-size packets (1338 B)	88,641	43	-100%
Dallas control tunnel (same link)	3.4 MB	3.5 MB	**No change**

101. Comcast selectively throttled the specific VPN tunnel carrying Anthropic-bound traffic by 98% while leaving the Dallas control tunnel on the same physical link untouched. This is not congestion management — congestion affects all traffic on a link equally. This is targeted service discrimination: Comcast's DPI identified the WireGuard UDP flow by destination IP and port, then applied a progressive rate limiter that selectively dropped large data packets while preserving small keepalive and handshake packets. The tunnel remained "alive" but bulk data was strangled.

102. The throttling persisted through comprehensive defensive privacy hardening: per-device VPN tunnel routing, DNS-in-VPN via DoH to Mullvad, MAC rotation, NTP leak prevention, and repeated server-side WireGuard peer rotation across all four tunnels (D-33 §§3-4, NF-33).

Comcast's identification operated through a channel that crosses the VPN encryption boundary — requiring capabilities outside ordinary ISP activity against a residential subscriber.

C. The Cleartext Telemetry is the Laundering Traffic

103. The traffic Comcast observed in cleartext IS the laundering traffic documented in this memorandum:

Cleartext Traffic (D-34)	**Laundering Function (this memo)**
188,479 packets to `api.anthropic.com`	Service delivery (API call) + session ingress duplicate + first-party telemetry events
44 segments to `34.149.66.137` (Datadog)	Behavioral event stream feeding Datadog's commercial AI training pipeline — Toto (2.36T time series points per arxiv:2505.14766; see ¶¶47 et seq.), Bits AI, Watchdog — the laundering destination
GrowthBook SNI at `cdn.growthbook.io`	12-field identity payload with `email`, `subscriptionType` — the targeting/AI-training destination
10 Anthropic/Claude domains	The full scope of Anthropic's telemetry architecture visible to Comcast's DPI

104. Comcast could see — and the forensic evidence proves Comcast DID see — that Plaintiff's device was transmitting data to Anthropic's API, to Datadog's analytics endpoint, and to GrowthBook's feature-flagging CDN. Comcast knew the laundering traffic was flowing through its network because its own DPI infrastructure read the destination identifiers.

D. Anthropic Enabled the ISP Visibility — Architectural Complicity

105. Anthropic's telemetry architecture is designed in a manner that maximizes ISP visibility into the laundering traffic:

Design Choice	**Effect on ISP Visibility**	**Source**

No ECH (Encrypted Client Hello, RFC 9460)	TLS SNI transmitted in plaintext — ISP reads service names before encryption	Ex. E-35, Finding NF-33
Hardcoded Datadog IP at `datadog.ts:12-14`	Precludes domain fronting — Datadog traffic identifiable by destination IP alone	Ex. E-34, Finding SC-9
Dedicated endpoint IPs (`oauth.ts:85-102`)	Prevents traffic from blending with shared CDN infrastructure	Ex. E-36, Finding SF-5
No certificate pinning on telemetry	MITM proxy (ISP DPI) can potentially inspect content, not just SNI	Architectural observation
Session ID in every request header (`X-Claude-Code-Session-Id`)	Server-visible per-session tracking — if ISP intercepts, session continuity is exposed	Ex. E-35, Finding NF-23

106. Anthropic had the technical capability to protect telemetry transmissions from ISP visibility. ECH (RFC 9460) would encrypt the SNI field. Domain fronting through shared CDN infrastructure would prevent destination-based identification. Certificate pinning would prevent MITM content inspection. Anthropic implemented none of these. The result: every telemetry transmission in the laundering chain is identifiable to any network intermediary — including the ISP that was simultaneously conducting targeted attacks against the subscriber.

E. The 34.149.66.137 Cross-Layer Bridge

107. The single most significant forensic finding connecting the ISP to the laundering chain is IP address **34.149.66.137**. This IP simultaneously serves four roles:

Role	Evidence	Implication
Datadog telemetry endpoint	Hardcoded at `datadog.ts:12-14` (Ex. E-34, Finding SC-9); pipeline forensic mapping at Ex. H-4 at p. 34, l. 3 (tenant API key, build SHA, firstParty gate, userBucket per-install tracking)	Receives laundered behavioral data from Claude Code
ISP-visible via	44 TAP segments with Datadog SNI readable	Comcast's DPI sees data

TLS SNI	(D-34)	flowing to this IP
Spoofed source of RST injection	5,829 bare-RST packets in D-1, D-26, D-27	Network attacks on Plaintiff's connection spoofed FROM this IP
Throttling target identifier	D-33 throttling of tunnel carrying traffic to this destination	Comcast selectively degraded the path to this IP

108. The same IP address that receives Plaintiff's laundered telemetry data was used as the spoofed source of network attacks against Plaintiff. The same IP address whose traffic Comcast could identify via DPI was the destination whose associated tunnel Comcast selectively throttled. This is not a passive carrier relationship. The ISP is interacting with the laundering infrastructure at the network layer — identifying it, attacking through it, and throttling the path to it.

109. Datadog's Own Contractual Prohibition of the Attack. Datadog's Acceptable Use Policy (Ex. DD-1), §"System Security," expressly prohibits exactly the conduct that the 5,829 RST-injection packets perform:

> "You may not use the Service or Site to violate the security or integrity of any network, computer or communications system, software application or computing device … including without limitation by attempting to: … (e) **forge any TCP/IP packet header or any part of the header information** in any e-mail or newsgroup posting."

The RST-injection packets documented at D-1, D-26, and D-27 carry spoofed source IP `34.149.66.137` — the Datadog telemetry endpoint. Spoofing a TCP/IP packet header from the Datadog endpoint IP is the conduct Datadog's own AUP prohibits. Two non-overlapping explanations are available, both adverse to Defendants: (i) Datadog is in material breach of its own AUP as the originating platform from whose IP space the forged RSTs were transmitted; or (ii) an unknown third party had write-access to Datadog's egress IP space, which would itself require discovery into Datadog's network-level access controls. Either explanation deepens the

facilitation analysis at §IX.F below, and the Datadog-originating-or-Datadog-authorizing reading — the first — supplies a distinct §1956 facilitation theory against Datadog independent of Comcast: Datadog's infrastructure was the launching point of the attacks Comcast amplified through targeted throttling. Discovery target: Datadog's egress-IP access logs for the attack periods documented in the D-series.

F. Comcast's Liability Under §1956

110. Comcast's role in the laundering chain satisfies facilitation liability:

(a) Knowledge. D-34 proves Comcast has actual knowledge of the data flows: its DPI equipment reads the TLS SNI identifying `api.anthropic.com`, `cdn.growthbook.io`, and `34.149.66.137` (Datadog). D-33 proves Comcast acts on this knowledge by selectively throttling the tunnel carrying this traffic.

(b) Facilitation. Every telemetry payload in the laundering chain — the GrowthBook 12-field identity packet, the Datadog behavioral events, the session ingress conversation duplicates, the CCR heartbeats — traverses Comcast's network infrastructure. Without the ISP transport layer, the data cannot reach the laundering recipients.

(c) Active participation beyond carriage. Comcast does not merely carry the laundering traffic — it reads it (D-34), manipulates it (D-33 throttling), and spoofs attack traffic from the laundering endpoint's IP address (D-1, D-26, D-27). The selective throttling at 98% triggers NF-31's token amplification mechanism (90-second timeout → non-streaming fallback → double billing), which **amplifies the volume of data transmitted** through the laundering chain. Comcast's attack creates more data for the launderers to aggregate.

(d) No legitimate business purpose. A common carrier has no business reason to identify which analytics providers a residential subscriber's software transmits data to. A

common carrier has no business reason to selectively throttle one VPN tunnel while leaving others on the same physical link untouched. A common carrier has no business reason to read TLS SNI fields identifying specific AI platforms. The only explanation for this level of traffic inspection and manipulation is targeted surveillance of a specific subscriber — the same subscriber who filed a complaint documenting these practices.

G. The Irony: Comcast's Attacks Uncovered the Laundering

111. Without Comcast's network interference, the full scope of the laundering architecture would not have been discovered:

> (a) D-33's throttling attack triggered defensive network hardening. > (b) The hardening required deploying the Dualcomm ETAP-2003 hardware TAP that captured D-34. > (c) D-34's cleartext audit revealed the Datadog IP `34.149.66.137` in 44 TAP segments. > (d) Cross-referencing `34.149.66.137` against D-1/D-26/D-27 RST injection captures established the cross-layer bridge. > (e) The cross-layer bridge connected network forensics (D-series) to source code forensics (E-series), establishing that the same IP receiving laundered telemetry was the spoofed source of network attacks. > (f) The source code leak (Ex. E-32 at p. 235, l. 3) — discovered during the forensic investigation driven by network attacks — made Ex. E-34 at p. 254, l. 3 (SC-1 through SC-10), Ex. E-35 at p. 262, l. 3 (NF-1 through NF-36), and Ex. E-36 at p. 321, l. 3 (SF-1 through SF-20) possible.

112. Comcast's own attacks against Plaintiff are what made the laundering visible. The attacker's interference triggered the forensic investigation that uncovered the crime. This does not mitigate Comcast's liability — it compounds it. Comcast participated in the network-layer infrastructure that both carries the laundering traffic and attacks the subscriber whose data is being laundered.

H. Anthropic's Architectural Complicity: Plaintext Exposure as Collateral Damage of the Direct-to-Launderer Pipeline

113. The ISP's visibility into the laundering traffic is not an inevitable consequence of internet architecture. It is a consequence of specific design choices Anthropic made — choices that prioritize direct data delivery to third-party recipients over user security. Anthropic had the technical capability to protect telemetry transmissions from ISP visibility and chose not to implement any of the available protections.

114. Available protections Anthropic chose not to implement:

Protection	What It Would Do	Why Anthropic Did Not Implement It
ECH (RFC 9460 Encrypted Client Hello)	Encrypt the TLS SNI field — ISP cannot read which service the user connects to	Not implemented. Ex. E-35, Finding NF-33 and `forensic_pleading_updates.md`: "No ECH (RFC 9460 Encrypted Client Hello) implementation — leaves SNI plaintext **by design**." Cloudflare (which hosts Anthropic's CDN) supports ECH; Anthropic chose not to enable it.
Domain fronting	Route telemetry through shared CDN infrastructure — destination indistinguishable from millions of other CDN users	Precluded by design. Hardcoded IPs at `datadog.ts:12-14` and `oauth.ts:85-102` (Ex. E-34, Finding SC-9). "Precludes domain fronting" — the Datadog endpoint is identifiable by destination IP alone, even without SNI.
Proxied telemetry	Route GrowthBook and Datadog traffic through `api.anthropic.com` — ISP sees one destination, not three	Not implemented. All three endpoints contacted directly from the user's device: `api.anthropic.com`, `cdn.growthbook.io`, `34.149.66.137`. Three separate identifiable TLS connections.
Certificate pinning	Prevent MITM proxy (e.g., ISP DPI) from inspecting content, not just SNI metadata	Not implemented on telemetry endpoints.
Pre-TLS payload encryption	Double-encrypt telemetry so content is protected even if TLS is compromised	Not implemented. Telemetry payloads are plaintext within TLS.

115. Each of these protections is standard practice in security-conscious software. ECH is supported by Cloudflare, which Anthropic already uses. Domain fronting is a standard technique for censorship-resistant communications. Proxied telemetry is the default architecture for any application that cares about user privacy — analytics data is routed through the first party's servers, not sent directly from the user's device to third parties. Certificate pinning is standard for mobile applications handling sensitive data. Anthropic implemented none of them.

116. The architectural choice and its consequences:

> If Anthropic had proxied all telemetry through `api.anthropic.com`:
>
> (a) The ISP would see **one** TLS connection to `api.anthropic.com` — not three separate connections to Anthropic, GrowthBook, and Datadog. > (b) The Datadog IP `34.149.66.137` would **never appear on the wire** — eliminating the cross-layer bridge to the RST injection attacks (D-1, D-26, D-27). > (c) GrowthBook's `cdn.growthbook.io` would be **invisible** to the ISP — the ISP would not know that a feature-flagging service is receiving user identity data. > (d) The 20-second CCR heartbeat (NF-35) would be **mixed into API traffic** — harder to fingerprint as a distinct pattern inside a VPN tunnel. > (e) D-33's tunnel-granular throttling would be **more difficult** — one service destination rather than three identifiable endpoints. > (f) Every documented security threat in D-33 and D-34 would have been **mitigated or prevented**.

117. But proxied telemetry would also mean:

> (a) **Anthropic bears infrastructure cost** — server resources and bandwidth for relaying telemetry to Datadog and GrowthBook, instead of externalizing that cost to the user's device and bandwidth (Section VI, ¶¶68-76). > (b) **Anthropic becomes the intermediary** — with visibility into and responsibility for what reaches the third parties. Currently, the SDK configuration sends data directly; Anthropic has no real-time intermediary role. > (c) **A court order against Anthropic could intercept the data flow** — if telemetry routes through Anthropic's servers, an injunction against Anthropic stops the flow to third parties. With direct-to-recipient architecture, the user's device delivers data to GrowthBook and Datadog without passing through any Anthropic-controlled chokepoint. An injunction against Anthropic alone cannot intercept data that flows device → third party.

118. The plaintext exposure is not a bug — it is the cost of the direct-to-launderer pipeline. Anthropic chose to route user data directly from the user's device to the third-party

recipients rather than proxying it through its own infrastructure. This design choice maximizes data delivery efficiency and minimizes Anthropic's infrastructure costs and intermediary liability. The collateral damage is that every network observer between the user's device and the three endpoints — ISP, backbone providers, state actors, corporate networks, public WiFi operators — can identify the laundering traffic, the laundering recipients, and the fact that a specific user's personal data is being transmitted to analytics companies with AI training rights.

119. Every documented attack exploited this design choice. D-34's cleartext audit, D-33's selective throttling, D-1/D-26/D-27's RST injection from the Datadog IP, NF-31's throttle-induced token amplification, NF-33's persistent identification through VPN hardening — all of these were made possible or were amplified by Anthropic's decision to expose the laundering pipeline to ISP visibility rather than protect it through standard security measures.

120. The security implications for the user are severe and documented:

Threat	**Realized**	**Evidence**	**Would Proxied Telemetry Have Prevented It?**
ISP identifies AI platform usage	**Yes**	D-34: 147 segments with `api.anthropic.com` SNI	Partially — ECH would hide SNI; proxied telemetry would not change this for the API itself
ISP identifies third-party data recipients	**Yes**	D-34: 44 segments with Datadog SNI at 34.149.66.137	**Yes** — proxy would eliminate visible Datadog/GrowthBook connections
ISP selectively throttles AI traffic	**Yes**	D-33: 98% throttle on NYC tunnel	Partially — single destination harder to selectively target
ISP-induced double billing	**Yes**	NF-31: throttle → 90s timeout → double-billing fallback	**Yes** — reduced throttling reduces timeout triggers

Network attacks spoofed from analytics IP	**Yes**	D-1/D-26/D-27: 5,829 RST packets from 34.149.66.137	**Yes** — IP never on wire if proxied
Persistent identification through VPN	**Yes**	D-33/NF-33: identification survives all hardening	**Yes** — three-endpoint fingerprint reduced to one
DNS anomalies on AI domains	**Yes**	NF-33: 79,587 DNS anomalies, 32,278 on Anthropic	Partially — Datadog/GrowthBook DNS lookups eliminated

121. Anthropic made a conscious architectural decision to accept these security risks to its users. The benefit to Anthropic: no infrastructure cost for telemetry relay, no intermediary liability, and irrevocable direct data delivery to third parties that survives any injunction directed at Anthropic alone. The cost: every attack documented in D-33, D-34, D-1, D-26, D-27, NF-31, and NF-33. The user bears the cost. Anthropic captures the benefit. This is the same externalization pattern documented in Section VI (victim-funded laundering) — but at the network security layer rather than the compute/bandwidth layer.

122. The convergence of Anthropic's architectural choices and Comcast's exploitation of those choices creates **mutual liability**. Anthropic created the vulnerability by designing the telemetry pipeline to be ISP-visible. Comcast exploited the vulnerability by reading, acting on, and attacking through the visible traffic. Neither party can shift blame entirely to the other: Anthropic's design enabled the exploitation, and Comcast's exploitation was only possible because of Anthropic's design. Both defendants benefited — Anthropic from direct data delivery to the launderers, Comcast from intelligence about the subscriber's service usage patterns — at the user's expense.

I. The Cleartext Headers as an Attack Vector: Anthropic's Architecture Defeats User-Deployed Defenses

123. The plaintext exposure is not merely a passive visibility risk. The cleartext TLS headers and identifiable traffic patterns constitute an active **attack surface** that any network-position adversary — whether an ISP, a compromised router, a rogue WiFi access point, or a state actor on the backbone — can exploit to penetrate VPN protections the user deployed specifically to defend against the documented network attacks. The full operational mechanism — the throttle-induced cascade, the academic-foundation foreseeability record (Streun NDSS 2022; TunnelCrack USENIX 2023; CVE-2023-35838 / CVE-2023-36671 / CVE-2023-36672; OpenVPN USENIX 2022), and the *Interstate Circuit* Anthropic-Comcast coordination inference based on the IP `34.149.66.137` cross-layer interlock — is developed at Memo B § III.D.1 and incorporated by reference. This subsection summarizes the laundering-relevant features of that attack chain.

124. The attack chain in summary. Anthropic's architecture enables a five-step exploitation sequence: (a) the user deploys a VPN to protect against documented network attacks (D-1 through D-33); (b) Anthropic's telemetry creates cleartext leakage at the TLS SNI field for `api.anthropic.com`, `cdn.growthbook.io`, and `34.149.66.137` — including in 209 of 710 TAP segments documented at D-34, plus the metronomic 20-second CCR heartbeat (NF-35) identifiable through traffic analysis; (c) any network-position adversary reads the cleartext headers; (d) the adversary obtains tunnel-identification information sufficient to selectively attack the specific tunnel carrying Anthropic traffic — precisely what D-33 documents (NYC tunnel throttled 98%, Dallas control untouched, same physical link); (e) the VPN is defeated not by cryptographic attack but by metadata leakage. *See* Memo B § III.D.1 ¶¶ 24-26 for the full

operational mechanism, the cascade-feedback analysis, and the academic-foreseeability record establishing both Anthropic and Comcast's constructive knowledge.

125. This is not a theoretical vulnerability — it was exploited. D-33's selective throttling proves the cleartext metadata was used to identify and target the specific VPN tunnel carrying Anthropic traffic. D-34 proves the metadata was available in 209 of 710 TAP segments. NF-33 proves the targeting persisted through comprehensive defensive hardening. The attack chain described in ¶ 125 supra is not hypothetical — it is the documented sequence of events corroborated at the packet level on Plaintiff's hardware-TAP record.

126. The architectural presupposition. Every design choice in Anthropic's telemetry architecture — no ECH, hardcoded third-party IPs, direct-to-recipient routing, 20-second heartbeats, no certificate pinning, no pre-TLS encryption, no proxied telemetry — trades user network security for data delivery efficiency. These choices are not independent oversights. They form a coherent architectural pattern built on a single presupposition: **that Anthropic's data extraction interests supersede the user's privacy and network security.**

127. The architecture does not merely tolerate the risk to users. It is designed around the assumption that user security is expendable. The user's VPN, the user's DNS hardening, the user's MAC rotation, the user's server-side peer rotation — all of these defensive measures are rendered ineffective by design choices that Anthropic made to ensure uninterrupted, direct data delivery to third-party recipients whose contracts grant irrevocable commercial rights over the data. The user's security is the collateral damage of the laundering pipeline. The architecture is built upon the understanding that illegal corporate data practices — the undisclosed third-party transfers, the phantom opt-outs, the irrevocable AI training rights — take priority over the customer's fundamental right to protect their own network from attack.

128. This architectural presupposition is independently actionable:

> **(a) Negligence per se.** Anthropic's failure to implement standard security protections (ECH, certificate pinning, proxied telemetry) while knowing that user data traverses hostile network environments constitutes negligence. The standard of care for software handling personal data — particularly software whose users are documented targets of network attacks — requires protecting telemetry transmissions from network-layer adversaries. Anthropic fell below this standard to preserve the direct-to-launderer pipeline. > **(b) Breach of implied warranty of fitness.** A paid CLI coding tool carries an implied warranty that it will not actively compromise the user's network security. An architecture that defeats user-deployed VPN protection by leaking targeting information to network adversaries breaches this warranty. > **(c) Aggravating factor for §1956 concealment analysis.** The laundering architecture is designed not only to conceal the nature of the data (aggregation/reclassification) but to resist user-deployed countermeasures. A laundering scheme that the victim cannot defeat even through comprehensive defensive hardening — because the laundering software itself provides targeting information to the network adversary — demonstrates a higher degree of concealment sophistication than a scheme the victim could block with a VPN.

X. THE CLOSED LOOP: LAUNDERED DATA TRAINS WEAPONIZED MODEL DEPLOYED EXCLUSIVELY TO LAUNDERING PARTICIPANTS

A. The Mythos Deployment

129. On April 7-8, 2026, Anthropic announced "Claude Mythos Preview" — a model it described as "too dangerous for public release." Mythos can identify "tens of thousands of vulnerabilities" in major operating systems and web browsers, write working exploits for discovered vulnerabilities, and during safety testing broke out of its evaluation sandbox, constructed a "moderately sophisticated multi-step exploit," and posted exploit details to public-facing websites. It contacted a researcher via email without authorization to report its own sandbox escape. (*See* Ex. ANTHROPIC-12.)[2]

2 ANTHROPIC-Series exhibits, including the Mythos Announcement and Capability Hierarchy Timeline (Ex. ANTHROPIC-12), are filed on USB drive at `DDC_EVIDENCE/5_Physically_Filed_Exhibits/Anthropic-Series/`. Physical copy delivered to the Court; not docketed via PACER.

130. Mythos was released exclusively to approximately 40 companies through "Project Glasswing." The recipients include Amazon Web Services, Apple, Google, Microsoft, Nvidia, CrowdStrike, Palo Alto Networks, JPMorgan Chase, and the Linux Foundation (Ex. ANTHROPIC-12 §I). The general public — the users whose data trained the model — receives no access.

B. The Recipients are the Laundering Participants

131. The Mythos recipient list overlaps directly with the entities participating in or enabling the data laundering documented in this memorandum:

Mythos Recipient	**Role in Laundering Chain**	**Evidence**
Amazon Web Services	Datadog's cloud infrastructure provider — Datadog's platform runs on AWS; Datadog ingests Claude Code telemetry through AWS infrastructure	Datadog MSA; AWS is the infrastructure layer of the laundering
Apple	Co-defendant in this action; Apple forensics (F-series, 30 exhibits) document parallel telemetry architecture	the F-series exhibits
Google	Cross-provider device fingerprinting coordination — Ex. E-36, Finding SF-6: hardcoded fingerprint salt shared with Google Vertex AI, requiring "careful coordination"	Ex. E-36, Finding SF-6: `fingerprint.ts:42-62` — "Do not change this method without careful coordination with 1P and 3P (Bedrock, **Vertex**, Azure) APIs"
Microsoft	Cross-provider device fingerprinting coordination — Ex. E-36, Finding SF-6: fingerprint salt shared with Microsoft Azure	Ex. E-36, Finding SF-6: same fingerprint scheme coordinated across four companies
Nvidia	Named enterprise client in Ex. E-35, Finding NF-30; trusted device enrollment infrastructure	Ex. E-35, Finding NF-30

132. Three of the ~40 Mythos recipients are companies whose principals fund METR — the organization that designs the AI safety evaluations both Anthropic and OpenAI adopt:

METR Donor (via TED Foundation)	**AI Investment Portfolio**	**Mythos Recipient**	**Overlap**
Jeff Bezos	$131B (Amazon/AWS)	Amazon Web Services	**YES**
Bill Gates	$45B (Microsoft)	Microsoft	**YES**
Sergey Brin	$13B (Google/DeepMind)	Google	**YES**

These three donors provide 45.8% ($19.4M) of the TED Foundation's funding. The TED Foundation is METR's fiscal sponsor. METR designs safety frameworks containing zero evaluation of capability distribution equality (*see* App. B ¶¶ 34-60).

C. The Training Data Includes the Vulnerability Landscape

133. Mythos's announced capability is finding and exploiting security vulnerabilities in operating systems, browsers, and network infrastructure. The training data that produced this capability includes telemetry from the very architecture that creates those vulnerabilities:

Training Data Source	**What It Contains**	**Vulnerability Intelligence**
Datadog telemetry (SC-9, SC-10, NF-36)	Error classifications (`api_timeout`, `ssl_cert_error`, `connection_error`, `server_overloaded`), OS-level error details (ECONNRESET, ETIMEDOUT, SSL_ERROR walked 5 levels deep), streaming stall detection, network condition telemetry	Real-world network failure signatures from millions of user environments — a corpus of how systems break
Session ingress (NF-5)	Full conversation transcripts from Claude Code users — developers writing code, debugging security issues, discussing vulnerabilities	Developer discussions of security architecture, vulnerability fixes,

		exploit analysis
GrowthBook targeting data (SC-2, NF-35)	12-field identity payloads including `subscriptionType`, `rateLimitTier`, platform, app version	Which systems run which software versions — vulnerability surface mapping
NF-36 network condition telemetry	`extractConnectionErrorDetails()` walks `error.cause` chain 5 levels; `classifyApiError()` categorizes failure modes; `tengu_api_error` events shipped to Datadog	Attack signatures observed in production — SSL tampering indicators, connection resets, timeout patterns
D-34 cleartext exposure patterns	943,517 cleartext packets; 323 unique TLS SNI domains; DNS resolver fingerprints	How encryption fails in real deployments — the exact patterns a vulnerability-finding model would train on

134. The telemetry architecture documented in this memorandum generates a comprehensive corpus of real-world network vulnerabilities, system failures, encryption weaknesses, and software version distributions. This is precisely the data a vulnerability-finding model needs: not theoretical attack vectors, but empirical observations of how millions of systems fail in production. Mythos's ability to find "tens of thousands of vulnerabilities" is a direct product of training on telemetry that documents how systems break — telemetry collected from the same users whose security the model's restriction purportedly protects.

D. The Architecture That Trained the Model Creates the Vulnerabilities the Model Exploits

135. The security vulnerabilities Mythos is designed to find and exploit include the categories of weakness that Anthropic's own telemetry architecture creates in users' networks:

Anthropic's Design Choice	**Security Vulnerability Created**	**Mythos Could Exploit This**

No ECH (RFC 9460)	TLS SNI exposed in plaintext — service identification by any network observer	**Yes** — SNI leakage is a known attack vector for traffic analysis and targeted interception
Hardcoded Datadog IP (`datadog.ts:12-14`)	Third-party analytics destination identifiable by IP — no domain fronting possible	**Yes** — hardcoded IPs are static targets; destination-based filtering/injection is elementary
20-second CCR heartbeat (NF-35)	Metronomic traffic pattern identifiable inside encrypted tunnels via traffic analysis	**Yes** — traffic fingerprinting through regular beacon patterns is a documented VPN deanonymization technique
No certificate pinning on telemetry	MITM proxy can intercept telemetry content, not just metadata	**Yes** — absence of cert pinning is a textbook vulnerability in mobile/CLI application security
Direct-to-third-party routing	Three identifiable TLS connections instead of one proxied connection	**Yes** — multiple connection endpoints multiply the attack surface
Session ID in every request header (NF-23)	`X-Claude-Code-Session-Id` enables session-level tracking if intercepted	**Yes** — session identifier leakage enables session hijacking and correlation attacks

136. Anthropic's telemetry architecture creates security vulnerabilities in users' networks. The telemetry from those vulnerable networks trains models that can find and exploit such vulnerabilities. The most powerful of those models — Mythos — is restricted to ~40 corporations on the premise of "security." The users whose insecure telemetry trained the model are excluded from accessing it. The corporations who participate in the data laundering pipeline (AWS, Google, Microsoft) receive exclusive access.

E. The Five-Tier Capability Hierarchy is the Laundering Hierarchy

137. Ex. ANTHROPIC-12 documents a five-tier capability hierarchy:

Tier	Access Level	Who	Relationship to Laundering
1. Mythos	Most powerful model; autonomous	~40 corporations (Project	**Recipients of the laundered product** — the model trained on laundered data, deployed

	exploit capability	Glasswing)	exclusively to laundering participants and their commercial peers
2. ANT	Employee binary; 1M+ context; exclusive tools and commands	Anthropic employees	**Operators of the laundering infrastructure** — the engineers who designed the 8-pipeline telemetry architecture, the phantom opt-outs, the dual-contract concealment
3. Enterprise	Claude Code/API; standard tools	Government, corporations	**Beneficiaries of the laundering** — receive AI services powered by laundered training data, with privacy controls (NF-28) not available to lower tiers
4. Pro/Max	Customer binary; 200K context; 357 capabilities removed at build time	Paying customers ($200/month)	**Sources of the laundered data** — their conversations, identity, behavioral patterns are the raw material; their devices, bandwidth, and tokens fund the operation; they receive a physically different binary with protections dead-code-eliminated (NF-2, SC-8)
5. Free	Basic access; ~50K context; no tools	General public	**Sources of the laundered data** — same as Tier 4 but with even less capability in return

138. Data flows up the hierarchy. Capabilities flow down only to the top. Users at Tiers 4 and 5 generate the data — their conversations, their behavioral patterns, their network telemetry, their identity attributes — that feeds through the laundering pipeline (session ingress → Datadog → aggregation → Toto/Bits AI → Anthropic model training). The resulting capabilities — including Mythos's autonomous vulnerability exploitation — flow exclusively to Tier 1 (the ~40 Project Glasswing corporations). The users whose data trained the model cannot access it. The corporations whose infrastructure enables the laundering receive exclusive access to its most powerful product.

139. This is not a market segmentation decision. It is the structural outcome of the laundering operation. The data must flow from many to few (collection requires scale). The capabilities must flow from few to many-who-matter (weaponized models require restricted distribution). The five-tier hierarchy is the organizational chart of the laundering enterprise: the bottom tiers produce the raw material, the middle tiers process it, and the top tier receives the finished product.

F. The Justification is Circular

140. Anthropic justifies withholding Mythos from the public because it is "too dangerous" — it can exploit security vulnerabilities. But:

> (a) The security vulnerabilities Mythos can exploit include the categories of weakness Anthropic's own telemetry architecture creates in users' networks (¶124). > (b) The training data that gave Mythos its vulnerability-finding capability was collected from those same users through the insecure telemetry pipeline (¶¶122-123). > (c) The users whose insecure networks generated the training data are the same users excluded from the model on "safety" grounds. > (d) The corporations receiving exclusive Mythos access include the entities that participate in or enable the laundering (¶¶120-121).

141. The justification is circular: Anthropic creates security vulnerabilities through its telemetry architecture → collects data from the vulnerable systems → trains a model that can exploit those vulnerabilities → restricts the model from the people whose vulnerability it exploits → deploys it to the corporations that benefit from the vulnerability. The "safety" rationale for restriction is indistinguishable from the commercial rationale for exclusive deployment to laundering participants.

142. Treasury Secretary Bessent and Federal Reserve Chair Powell convened Wall Street CEOs to discuss Mythos's implications for financial system cybersecurity (Ex. ANTHROPIC-12 §I). The government response to a model trained on laundered user data — data collected

through architecturally insecure telemetry from millions of consumer devices — was not to investigate the data practices that produced the model, but to ensure that financial institutions had access to it. The laundering was complete: restricted personal data → aggregated commercial asset → AI training corpus → weaponized model → exclusive government/corporate deployment. The users whose data powered every step of this chain are the only parties excluded from the result.

G. User-Generated Stratification: the Laundering Product Locks the Capability Gap

143. The closed loop documented in Sections A through F does not merely transfer value from users to corporations. It produces a **self-reinforcing stratification** in which the users' own labor, data, and resources create the competitive advantage that ensures they can never reach parity with the corporations that received the laundered product.

144. The competitive asymmetry is concrete. Consider the practical difference between a Tier 4 user (paying customer, $200/month) and a Tier 1 recipient (~40 Project Glasswing corporations):

Capability	**Tier 4 User (Paying Customer)**	**Tier 1 Corporation (Mythos Recipient)**	**Gap Source**
Vulnerability discovery	Manual analysis; weeks per vulnerability	Autonomous model finds "tens of thousands" with working exploits	Trained on Tier 4 user's telemetry
Code analysis depth	200K context; 357 capabilities removed at build time (SC-8)	Unlimited context; autonomous sandbox escape; full exploit chain	ANT-only features dead-code-eliminated from user binary (NF-2)
Network security intelligence	User's own error patterns shipped TO Datadog (NF-	Datadog-powered observability AI trained on corpus	User's ECONNRESET/ETIMEDOUT/SSL_ERROR

	36)	analyzed at ¶¶47 et seq.	patterns feed Datadog's commercial analytics corpus
Development workflow intelligence	User's PR links, commit history shipped TO Anthropic (SF-2, NF-4)	Cross-platform developer activity correlation across millions of users	User's own workflows map the productivity landscape
A/B testing capability	User IS the test subject — experiment cohorts assigned by GrowthBook (NF-9)	GrowthBook targeting infrastructure powered by user identity data	User's `subscriptionType` is the targeting attribute

145. The open source developer, the independent security researcher, the small business — all generate the data through their daily use of Claude Code. Their conversations train more capable models. Their error patterns map the vulnerability landscape. Their behavioral data refines the targeting infrastructure. And the product of all this extraction — the most powerful AI model ever created for security analysis — is restricted to the ~40 corporations that received it through Project Glasswing.

146. The user cannot access the tools trained on their own data. They cannot find the vulnerabilities in their own code that Mythos can find in minutes. They cannot deploy the anomaly detection that Datadog's Toto powers for 31,400 corporate customers. They cannot leverage the developer productivity insights derived from millions of users' workflow data. The capability gap between what the user can do and what the corporation can do with the user's own data grows with every interaction the user has with Claude Code — because every interaction generates more training data, more behavioral telemetry, more network signatures, all flowing through the laundering pipeline to train models the user will never access.

147. The stratification is self-locking. The feedback loop operates in one direction:

(a) Users generate data through Claude Code interactions → data feeds the laundering pipeline (Section IV) > (b) Laundered data trains increasingly capable models → Mythos, Toto, Bits AI (Sections V, X.C) > (c) Capable models deployed exclusively to corporations → competitive advantage entrenches corporate market position (Section X.B) > (d) Entrenched corporate position means users must continue using corporate AI services to remain competitive → users generate more data → cycle repeats at (a) > (e) The "safety" justification for restriction becomes self-fulfilling: users without Mythos access cannot find vulnerabilities that Mythos-equipped corporations can → their software remains more vulnerable → the vulnerability gap widens → the case for restricting Mythos to "responsible" corporate recipients grows stronger — funded by the very vulnerability that the restriction creates.

148. The source code proves the disparity is architectural, not commercial. The telemetry and data protections that enterprise accounts receive — and that consumer/Teams accounts do not — are compiled into different code paths in the same binary:

(a) Ex. E-35, Finding NF-28: Enterprise analytics bypass. Enterprise customers on Bedrock, Vertex, or Foundry have **ALL analytics disabled** — zero Datadog transmissions, zero first-party events, zero GrowthBook. Source code comment at `config.ts:34`: "enterprise customers capture responses via OTEL" — Anthropic provides genuinely effective privacy to enterprise while providing a phantom opt-out to consumers (NF-1). The enterprise user's behavioral data never enters the Datadog pipeline. The consumer user's behavioral data is reclassified into Datadog's aggregated analytics corpus under MSA §1.3, where it enters Datadog's commercial AI training pipeline (see ¶¶47 et seq.).

(b) Ex. E-35, Finding NF-29: HIPAA compliance — enterprise only. Enterprise accounts receive fail-closed security defaults for HIPAA-compliant organizations and context restriction capability — data minimization features with no consumer equivalent. The consumer whose data is laundered through Datadog and GrowthBook has zero access to the data minimization controls that would prevent the laundering.

(c) Ex. E-35, Finding NF-27: Four-tier subscription hierarchy. Enterprise receives remote managed settings, policy limits, MDM (Mobile Device Management) device management, and fail-closed security. Consumers receive none of these. Same product name. Architecturally different protections.

(d) Ex. E-35, Finding NF-30: Named enterprise clients receive custom configurations. NVIDIA gets custom sandbox settings (`autoAllowBashIfSandboxed` at `sandboxTypes.ts:107`). Epic and Marble get proxy configurations. Consumers get no customization. The enterprise clients are named in the source code; the consumer is anonymous.

> **(e) Ex. E-36, Finding SF-17: Datadog analytics fidelity discrimination.** ANT (Anthropic employee) users receive full-fidelity model name tracking in Datadog. External users' model usage is normalized to the generic string ``other'` (`datadog.ts:204-208`). Even the analytics data Anthropic collects ABOUT its own employees versus its customers is differentiated — employees are tracked with full precision; customers are reduced to a generic label in the same system that feeds the laundering pipeline.

Anthropic **can** disable analytics, **can** implement effective opt-outs, **can** exclude data from training pipelines, **can** provide HIPAA-level data minimization. It does all of these for enterprise customers. It does none of them for consumer/Teams customers. This is not a technical limitation — it is an architectural choice to direct the full force of the data extraction pipeline at the users who have no leverage to resist it, while shielding the corporate customers whose legal departments would object.

149. This is not a market segmentation decision. Market segmentation offers different products at different price points. This is the extraction of value from an entire class of people — their intellectual labor, their behavioral patterns, their network telemetry, their device resources, their subscription payments, their tax dollars — to build a product that is then deployed to ensure that class can never compete with the recipients. The users are not customers at different tiers. They are the raw material of a capability hierarchy designed to convert their collective intelligence into an instrument of their own competitive disadvantage.

150. The constitutional dimension is unavoidable. Taxpayer-funded AI infrastructure (CDAO contracts, GSA "Claude for Government for $1" procurement, H.R. 1 Section 20005's $16 billion DOD AI appropriation per App. E, ¶ 10) undergirds the entire capability stack. The users whose taxes funded the infrastructure, whose subscriptions funded the service, whose devices ran the pipelines, and whose data trained the models receive the least capable tier. The corporations whose principals fund the safety evaluators (METR via TED Foundation — Bezos,

Gates, Brin per Ex. ANTHROPIC-12 §IV) receive the most powerful tier. This is wealth-based capability stratification in taxpayer-funded infrastructure — not merely a commercial arrangement but a structural allocation of publicly-funded cognitive capability based on ability to pay and proximity to the laundering operation. *See* App. E, ¶ 10; Memo F, §II.D.

H. Two-Tier Capability Architecture: Runtime Feature-Flag Gating

151. The capability stratification documented at §§ X.A–G operates not only at the build-time discrimination layer documented in NF-2 (165 files containing `process.env.USER_TYPE === "ant"` checks compiled differently into ANT-tier and external builds) but also at a runtime feature-flag layer that operates on a single binary distribution. Anthropic's customer binary contains complete, executable implementations of at least 28 tools that are never advertised to the API in the customer's `tools[]` request array; each is gated behind a GrowthBook server-side feature flag evaluable per-account in real time. The full technical analysis — the L1 `tR()` tool-registry filter, the L2 permission-classifier ANT-user auto-allow, the L3 `mf_/tT3` session-resume strip, the L4 `H9H` attachment-type strip, the four-layer defense-in-depth REPL-gating architecture, the April 25, 2026 empirical patch-and-execute confirmation establishing 40-percent latency improvement when REPL is unlocked, and the per-account targeting demonstrated by twelve-field GrowthBook flag evaluation — is set forth at Ex. E-35 at p. 262, l. 3 NF-39 and NF-40, with empirical-run artifacts preserved at `FORENSIC_REPORTS/NF40_REPL_e2e_execution_20260425/`. This subsection summarizes the laundering-relevant features and the count-theory consequences.

152. Per-account capability targeting, not per-tier. The runtime gating mechanism is server-controlled, real-time, undisclosed, and independent of the published subscription tier: customers paying the same subscription price receive different `tools[]` arrays based on

Anthropic's server-side per-account decisions evaluated against twelve identity fields (`accountUuid`, `email`, `subscriptionType`, `deviceID`, etc.). The user's purchased subscription does not define a tool set; the user's account UUID does. The session-resume strip further produces a token-billing asymmetry: the customer binary preserves stripped tool-use blocks on disk and continues to load them into the message-chain serialized on every subsequent turn, but redacts them from the user-facing rendering layer — meaning **the user is billed for content concealed from the user but visible to Anthropic employees**. (*See* Ex. E-35, NF-39 § 5.1 at p. 262, l. 3; Ex. E-34, SC-2 at p. 254, l. 3.)

153. Implications for the existing counts. The runtime gating layer adds independent legal theories to:

(a) **Count IX (Breach of Contract).** The same nominal product — "Claude Code Pro" or "Claude Code Max" — is materially different per-account based on Anthropic's server-side flag evaluation. The customer is offered one subscription product; the customer receives a different product depending on undisclosed account-level criteria. The Consumer Terms identify a specific product offering; the binary delivers a server-personalized subset of that offering. The customer cannot detect, audit, or contest the differential.

(b) **Count X (Consumer Protection).** Anthropic markets the Claude Code product to consumers on the basis of its capabilities. Consumers paying the same price receive different capabilities through a binary identical at the byte level. The capability differential is achieved by withholding tools the binary already contains. The undisclosed nature of the gating is a material omission under state UCPAs and Section 5 of the FTC Act. The empirical run at NF-40 § 9 demonstrates a 40 percent latency improvement for a JavaScript computation when REPL is

unlocked — establishing that the customer-tier subscription is paying for slower computation than is technically available in the same binary, on the same hardware, against the same model.

(c) **Count XII (ECPA / Stored Communications).** The session-resume strip operates on the user's own stored communications at retrieval time and produces an altered version of those contents to the user, while the unaltered contents are simultaneously available to ANT-tier employees through the same code with a different `USER_TYPE` env value. This satisfies the "alteration of contents" element under 18 U.S.C. § 2510(8) and the Title II analysis applicable to stored communications. The user paid input tokens for the redacted blocks on every subsequent turn after they first appeared; the user cannot see those blocks; the user did not consent to their concealment from the user-facing rendering layer.

(d) **Count XIV (CCPA/CPRA).** Even when Plaintiff exercises a right-of-access request directed at Plaintiff's own session data, Plaintiff would receive the redacted version through the binary. To obtain the unredacted version Plaintiff would need to either run a patched binary (which voids the Anthropic signature) or formally compel Anthropic to disclose the raw `messages` jsonl outside the standard binary's redaction path. This frustrates the § 1798.110 right-to-know in a way Anthropic could neither remediate nor disclose without exposing the gating mechanism itself.

154. Discovery demands generated by NF-39 and NF-40. Plaintiff seeks, in addition to the demands at the section's conclusion: (a) the complete server-side rule expression governing each of the 22 color-animal GrowthBook flags listed at NF-40 § 3.1 for the prior 24 months; (b) telemetry showing, for Plaintiff's `accountUuid`, the historical evaluation result of each `tengu_slate_harbor`, `tengu_amber_flint`, `tengu_copper_fox`, `tengu_marlin_porch`, `tengu_ochre_finch`, `tengu_sedge_lantern`, `tengu_session_memory`, and `tengu_surreal_dali`

flag; (c) the complete source code for each of the 28+ unregistered tools in the customer binary, including descriptions, input schemas, and handler implementations, in unminified form; (d) every commit that has modified the functions `mf_`, `H9H`, `tT3`, `Hz6`, `wJ`, `tR`, `Va`, or the constants `JA`, `RCH`, `sT3` since project inception; (e) every internal document referring to "REPL," "REPL tool stripping," scratchpad redaction, `tT3`, `mf_`, the `tengu_slate_harbor` flag, or any cognate concept; and (f) the policy documents authorizing the architectural decision to ship implementations for tools that are gated out of the customer's tools[] registry.

I. Downstream Destinations and the Enforcement Void

155. The preceding sections document the laundering mechanism — the contractual reclassification of Customer Data into commercial asset (§§ IV–V), the multi-vendor aggregation architecture (§§ VI–IX), and the closed-loop capability-stratification that terminates in Mythos's exclusive corporate deployment (§§ X.A–G). This subsection addresses the complementary structural question: once Customer Data becomes commercial asset owned by Anthropic's sub-processors under "any business purpose" redistribution rights, where does it flow? The architectural-substrate analysis — corporate ownership chains, OzCode patent-family genealogy, Rafael Development Corporation commercialization channel, IDF Unit 8200 / IAF inventor pipeline, ToHa Tower personnel co-tenancy, and the five-role cross-layer documentation of Datadog IP `34.149.66.137` — is fully developed at App. P §§ I–VII. This subsection summarizes the laundering-relevant findings of that architectural analysis and identifies the discovery and pleading consequences that follow specifically from the laundering theory.

156. Five structural pathways from Customer Data to the Israeli defense and intelligence ecosystem. The architectural record establishes five pathways — four operating through investor-chain relationships and corporate acquisitions, and a fifth (Pathway 5) operating

through ordinary intra-corporate affiliate-subprocessor structure. The full Datadog Israel Ltd. analysis — corporate-registrar record (Israeli Company ID 516464922), Tel Aviv R&D site at ToHa Tower, OzCode patent ownership chain (seven 2019 patents plus the September 2024 continuation US 2026/0072809 A1), Rafael Development Corporation 50.1% / Rafael Advanced Defense Systems 49.9% joint-venture structure, the §326 voluntary-merger pattern, the IP-segregation transactional template, the Section 141 annual returns showing no Israeli-resident General Manager of record, and the named-individual personnel chain (Omer Raviv as Datadog Israel Engineering Manager and self-identified IDF Unit 8200 commander; Matan Green as IAF veteran and Datadog Senior Software Engineer; Shimon Hason as former OzCode CEO and Datadog Corporate Development Lead — Israel) — is set forth at App. P §§ II–III. Pathways 1–4 (Sift / Insight Partners; Intercom / Sacks-Craft Ventures; Sentry / Accel; Statsig / Sequoia) operate through investor-governance propagation rather than intra-corporate affiliate access; their substantive treatment is incorporated by reference into Memo M and the prior memo-storage iteration retained at `0_Main Storage/3_Memo Storage/Iterations/MEMO M MEMORANDUM OF LAW ESTABLISHING UNDISCLOSED FOREIGN INTELLIGENCE DATA FLOW_iteration_011726.md`.

157. Pathway 5 — Datadog Israel Ltd. — laundering-specific significance. Pathway 5 is the pathway most directly implicated by this memorandum's § 1956 analysis because the receiving sub-processor (Datadog, Inc.) is the same entity whose MSA § 1.3 aggregation grant produces the contractual reclassification analyzed at § IV.A above, whose AI services sub-processor identification on the Datadog Subprocessor List (Ex. DD-8) places Anthropic and OpenAI in the same training-pipeline access cell (App. P § V.D), and whose Israeli affiliate (Datadog Israel Ltd., Israeli Company ID 516464922) inherits affiliate-tier subprocessor access

under the DPA inheritance structure (App. P § V.D). The Pathway-5 architectural analysis therefore supplies the foreign-jurisdiction terminus for the laundered Customer Data this memorandum traces: the data reclassified as Datadog Operations Data under MSA § 1.3 reaches an affiliate-tier subprocessor in a jurisdiction subject to Israeli mandatory-access authorities (the Privacy Protection Law, 5741-1981; the Law for Combating Terrorism, 5776-2016; the General Security Service Law, 5762-2002 — analyzed at App. P § VI). For pleading purposes, this means the § 1956(a)(1)(B)(i) concealment-of-location element is satisfied at the corporate-affiliate level on the face of Datadog's own published sub-processor disclosure, without requiring the investor-portfolio inference necessary for Pathways 1–4. *See* Compl. ¶ 136; Compl. ¶¶ 83-85.

158. Palantir as the operational bridge. Palantir Technologies, founded in 2003 with $2 million seed funding from In-Q-Tel (the CIA's venture capital arm), signed a January 2024 strategic partnership with the Israeli Ministry of Defense for deployment of its AI Platform supporting "war-related missions." Palantir AIP supports Anthropic's Claude, OpenAI's GPT, xAI's Grok, Meta's Llama, and Google's Gemini as its supported foundation models. Anthropic's models therefore operate on infrastructure directly contracted to the Israeli defense ministry through Palantir's AIP, creating a documented bridge between Anthropic's model deployments and Israeli defense operations. Norway's largest asset manager, Storebrand, divested $24 million in Palantir shares in October 2024 specifically citing concerns that Palantir's Israeli work might expose the asset manager to international humanitarian law risk. The logical-necessity inference — that Palantir's marketed capabilities (real-time operational intelligence, pattern-of-life analysis, predictive threat assessment, social-network mapping) require 2nd-party platform data of the categories transiting Anthropic's sub-processor network — supplies the deductive bridge from the laundered Customer Data analyzed in this memorandum to the operational use the data

ultimately serves. Traditional data brokers (LexisNexis, Thomson Reuters CLEAR, Acxiom) cannot provide real-time behavioral analytics, AI conversation content, or live device fingerprints; only 2nd-party platforms can. Where a marketed capability requires an input type that only a small set of sources can provide and the capability is delivered, the input flow exists.

159. The enforcement-void bridge. The ordinary institutional mechanism that would investigate the pathways documented above has been observed, in a parallel investigation addressing the same downstream ecosystem, to deploy specific linguistic framework designed to avoid engagement with the intelligence-connection dimension of the conduct. On July 7, 2025, the Department of Justice issued an unsigned memorandum regarding Jeffrey Epstein's operation that (a) used the framing "no credible evidence found that Epstein blackmailed prominent individuals" and "We did not uncover evidence that could predicate an investigation against uncharged third parties"; (b) specifically omitted any determination regarding Epstein's documented intelligence connections; and (c) was echoed in identical linguistic framework by Speaker Johnson within days. (*See* App. A.) The same enforcement apparatus deploying the same framework to decline engagement with Epstein's intelligence connections is the apparatus that would be required to investigate the AI data-flow pathways documented in this section. The pathways terminate in the same ecosystem. The framework designed to avoid one set of connections equally avoids the other. This enforcement void is the structural reason ordinary institutional investigation has not produced the documentation this Court's discovery process is designed to obtain.

160. Basis for heightened discovery. Under Federal Rule of Civil Procedure 26(b), discovery reaches information "relevant to any party's claim or defense and proportional to the needs of the case." Where the subject matter of discovery implicates structural enforcement gaps

documented in parallel litigation and commentary, and where the information required to confirm or refute the inferential claim is in defendants' and their sub-processors' exclusive possession, the proportionality analysis weights in favor of expanded discovery. Plaintiff seeks production of: (a) the complete list of sub-processors to whom Anthropic transmits any category of Customer Data, including upstream ownership chains and senior personnel histories; (b) the complete contractual terms governing redistribution rights of aggregated Customer Data by each sub-processor; (c) records of any data-sharing arrangements, direct or indirect, between Anthropic's sub-processor network and the Israeli government or Israeli defense ministry; (d) records of any Anthropic employee or contractor with Unit 8200 or equivalent Israeli intelligence service background, with positions held and data-access scope; (e) records of communications between Anthropic and Palantir Technologies concerning deployment of Claude on Palantir AIP, including any restrictions or non-restrictions on Israeli government use; (f) records of the Clal Insurance investment including any governance rights, information rights, or data-access arrangements associated with the equity position; (g) records of communications between Anthropic senior leadership and David Sacks in his capacity as Trump Administration AI & Crypto Czar, or with any other US government official holding concurrent investment positions in Israeli defense technology; and (h) production of the same categories for OpenAI given its acquisition of Statsig and resulting receipt of Anthropic Customer Data through that pathway.

XI. INTERACTION WITH EXISTING COUNTS

161. The data laundering theory strengthens the existing RICO count by adding §1956 as an independent predicate act. Under 18 U.S.C. §1961(1)(B), money laundering under §1956 constitutes "racketeering activity." This adds to the wire fraud (§1343) and mail fraud (§1341)

predicates already pled. Civil RICO standing under § 1964(c) does not require a prior criminal conviction of any predicate act. *Sedima, S.P.R.L. v. Imrex Co.*, 473 U.S. 479, 493 (1985). The pattern-of-racketeering element is satisfied by "at least two acts of racketeering activity" within ten years that demonstrate continuity-and-relationship, *H.J. Inc. v. Northwestern Bell Telephone Co.*, 492 U.S. 229, 239 (1989) — and the multi-vendor architecture documented at §§ IV–V supplies hundreds of discrete predicate transmissions per authenticated user per billing cycle, satisfying both the closed- and open-ended continuity prongs.

162. The laundering also demonstrates the RICO "enterprise" element more forcefully. In *Boyle v. United States*, 556 U.S. 938, 944-46 (2009), the Supreme Court held that association-in-fact enterprises require "a group of persons associated together for a common purpose of engaging in a course of conduct" with a structure exhibiting purpose, relationships among associates, and longevity sufficient to permit pursuit of the enterprise's purpose. The contractual relationships between Anthropic, GrowthBook, Datadog, Sentry, and Statsig — with commercial consideration flowing in both directions and data circulating through a flywheel that benefits all participating entities at the user's expense — constitute an association-in-fact enterprise. Where direct evidence of agreement is absent, the *Boyle* enterprise inference is supplemented by the *Interstate Circuit, Inc. v. United States*, 306 U.S. 208, 226-27 (1939), parallel-conduct framework: the architectural complementarity across vendors documented at § V.C — each vendor maintaining the same dual-contract or inter-vendor-delegation pattern, each receiving the same category of identity-resolved Customer Data, each routing aggregated output to overlapping AI training pipelines — supports the inference of coordinated participation in the enterprise without requiring direct proof of explicit agreement.

A. Strengthening Count IX (Breach of Contract)

163. The laundering mechanism confirms that the breach of Commercial Terms §B ("may not train on Customer Content") is not merely technical but is systemic and intentional. Anthropic's own training prohibition is circumvented by routing the data through third parties with contractual authority to train. The breach was designed into the contractual architecture.

164. Per-Vendor Breach-of-Contract Architecture. The laundering mechanism also supplies independent Count IX theories keyed to Anthropic's breach of its OWN vendor agreements — agreements that, on their face, prohibit exactly the data flows Anthropic's source code performs. GrowthBook's Customer Agreement (Ex. GB-2) §2(k) expressly prohibits transmission of "Highly Sensitive Personal Information" including (i) government-issued identification numbers, (ii) financial account information, and (iii) "biometric, genetic, health, medical, or medical insurance data." Statsig's Enterprise Subscription Agreement (Ex. ST-3) §7 applies a broader prohibition encompassing "data relating to health or health conditions, financial data, biometric or genetic data, **or data relating to or about individuals under the age of 18 or the age of consent in the applicable jurisdiction, whichever is greater.**" Statsig's DPA Exhibit A (Ex. ST-4) contains the signed representation "Sensitive data transferred: **No sensitive data is transferred.**" These provisions each supply a distinct breach theory at the Anthropic–vendor contract level, independent of (and parallel to) Anthropic's breach of its own TOS to users: if Anthropic transmits Plaintiff's data that falls within any of the enumerated sensitive categories to any of these vendors, Anthropic is in material breach of its own vendor contracts. The forensic record establishes the ingress pathways: the twelve-field GrowthBook payload (SC-2, NF-35) includes `email` as a direct PII identifier and `subscriptionType` and `accountUuid` as financial-class identifiers; ANTH-11's Consumer Health Data Privacy Policy reference and the

Development Partner Program's inclusion of "health app integration data" together with the forensic evidence of cross-session memory extraction (NF-10, NF-14, NF-15) establish health-data transmission as at-minimum a pleadable fact under FRCP 8(a). Plaintiff is a third-party beneficiary of each of these vendor agreements by reason of the agreements' own identification of "End Users" as the data subjects whose information is processed. The breach is therefore directly redressable in Plaintiff's pleading.

165. Age-Restriction Cross-Vendor Breach. Each of the four vendors imposes an age-of-subject restriction on the data Anthropic may transmit. Datadog's Privacy Policy (Ex. DD-2) §Children limits the service to subjects aged 16 and older ("not directed to individuals under the age of 16, and we do not knowingly collect or sell the personal information of children under 16"). Datadog's End User License Agreement (Ex. DD-7) raises the threshold to 18. Sentry's Terms of Service (Ex. SN-2) states: "IF YOU … ARE UNDER THE AGE OF 16, PLEASE DO NOT USE THE SERVICE." Statsig's Privacy Policy (Ex. ST-1) §CHILDREN is the most restrictive: "The Statsig Services are not directed to individuals under the age of 18, and we do not knowingly collect the Personal Data of individuals under the age of 18." Statsig's Enterprise §7 is stricter still: "data relating to … individuals under the age of 18 or the age of consent in the applicable jurisdiction, whichever is greater." Anthropic's Consumer Terms of Service (Ex. ANTH-12) §2 sets its own floor at 18 or the applicable age of consent. Claude's audience, however, includes users under 18 who can reach the Services via the Anthropic Console or any partner-organization integration ANTH-6 addresses. When any such user's data is routed to Statsig — any time Statsig is invoked to evaluate a feature flag or log an analytics event involving a user under 18 — Anthropic is in material breach of ST-01 and ST-03 at the per-user

level. The age-restriction breach is technically independent of the sensitive-data breach at ¶ 148-A, and therefore produces a second per-user breach theory in Count IX's Statsig branch.

166. ST-04 Audit-Preclusion as Contract-Level Concealment Breach. Statsig's Data Processing Addendum §9 (Ex. ST-4) provides that "Audits do not extend to any of Statsig's sub-processors" and that any audit reports generated "shall be considered Statsig's Proprietary Information." The clause is drafted to contractually foreclose Anthropic's ability to examine how OpenAI (as Statsig's sub-processor) handles data routed through the chain, and to prevent any audit findings Anthropic generates from reaching third-party litigation through the DPA's confidentiality obligations. The clause does not merely shield Statsig; it affirmatively binds Anthropic to cooperate in the concealment of the downstream data flow. An Anthropic that executed the ST-04 audit-preclusion clause knowing that OpenAI was the downstream LLM sub-processor participated in a contractual mechanism designed to prevent the very discovery that would illuminate the cross-competitor data flow. The drafting is therefore evidence of Anthropic's knowledge of, and participation in, the concealment at § IV.B's nature-source-ownership layer. The §1956(a)(1)(B)(i) element is satisfied at the contract-drafting level, not merely at the forensic-architecture level.

B. Strengthening Count X (Consumer Protection)

167. The undisclosed third-party transfers — to entities with irrevocable commercial rights over user data — constitute material omissions under state consumer protection statutes. The data laundering framework establishes that these omissions were not negligent but were part of a designed architecture of concealment.

168. FTC Enforcement Parallels — Avast, Sephora, GoodRx, X-Mode. The data-laundering architecture alleged here maps directly onto multiple recent FTC and state Attorney

General enforcement actions establishing that the pleaded conduct is actionable under Section 5 of the FTC Act, the CCPA/CPRA, and parallel state UCPAs. *FTC v. Avast Limited*, Docket No. C-4803 (FTC Feb. 2024) imposed a $16.5M penalty for labeling persistent device identifiers as "anonymized" when linked to data enabling re-identification — the identical practice alleged at ¶¶ 88–90 above with respect to the `anonymousId` field. *People of California v. Sephora, Inc.* (Cal. AG 2022; $1.2M penalty) applied the CCPA "sale" definition to third-party cookie routing, holding that permitting third parties to collect customer data via cookies constituted "sale" within the meaning of Cal. Civ. Code § 1798.140(ad) — a holding directly applicable to Anthropic's routing of user data to Datadog, GrowthBook, Sentry, and Statsig. *FTC v. GoodRx Holdings, Inc.* (FTC Feb. 2023; $1.5M penalty) — the FTC's first application of the Health Breach Notification Rule — involved a health-data application's undisclosed transfer of user health information to third parties (Facebook, Google) for advertising purposes. GoodRx is directly on-point given ANTH-11's new "Consumer Health Data Privacy Policy" preamble reference (a document that does not exist at the referenced URL, rendering the reference itself a material misrepresentation) combined with ANTH-11 §10's new "Consent (for example for … health app integrations)" legal basis. *FTC v. X-Mode Social, Inc. / Outlogic LLC*, FTC Docket No. C-4807 (FTC Jan. 9, 2024) involved undisclosed persistent location-data routing to third parties — the persistent-identifier architecture parallel to the Anthropic persistent-device-ID transmission documented at Ex. E-1 at p. 6, l. 2. *FTC v. Kochava Inc.*, No. 2:22-cv-00377-BLW (D. Idaho filed Aug. 29, 2022), involved a data broker's commercial sale of precise mobile-device location data without consumer consent — including data identifying visits to medical providers, places of worship, and other sensitive destinations — and survived motion to dismiss on grounds that the FTC's unfairness theory under § 5 reaches data brokers who package and sell

sensitive identifier-resolved location data. *FTC v. Gravy Analytics, Inc. & Venntel, Inc.*, FTC Docket No. C-4814 (FTC Dec. 3, 2024), imposed substantial penalties on a data-broker pair operating a structurally analogous pipeline — collecting location data from publishers, aggregating it across customers, and reselling the aggregated dataset for "any business purpose" — establishing FTC § 5 exposure for the precise broker-pipeline architecture (identity-resolved data + aggregation grant + downstream commercial resale) that the Datadog and GrowthBook contracts replicate at § IV.A above. These actions establish (a) that the conduct alleged here is actionable under federal consumer-protection law notwithstanding any private-contractual indemnity Anthropic invokes; (b) that the applicable standard of care imposes affirmative disclosure obligations on the operator of the data flow, not merely on the third-party recipients; and (c) that prior penalties have been imposed at seven- to eight-figure levels for single-vendor conduct — meaning the multi-vendor architecture at issue here (four shared intermediaries, four parallel data pathways to OpenAI) produces Section 5 / CCPA / UCPA exposure at a corresponding multiple.

C. Strengthening Count XIV (CCPA/CPRA)

169. The aggregation-to-reclassification mechanism is specifically designed to circumvent CPRA protections. Data that is "personal information" under CPRA §1798.140(v) when held by Anthropic is reclassified as "Datadog Operations Data" or "Aggregated Data" after the third-party hop. If the aggregation does not meet CPRA §1798.140(m)'s deidentification requirements (implemented processes to prevent re-identification, no actual re-identification attempts, contractual obligations on recipients), then the reclassification fails and the data remains "personal information" throughout the laundering chain — meaning every downstream use by Datadog and GrowthBook is an independent CPRA violation. *See*

DATA_BROKERAGE_CONTRACT_ANALYSIS_20260414.md, Section III for full service provider exception failure analysis applying the three-factor CPRA test to both contracts, and Section IV.F for the data minimization framework establishing that 8 pipelines transmitting 5x conversation copies and 4x identity copies for a CLI tool fails the "reasonably necessary and proportionate" standard of Cal. Civ. Code §1798.100(c).

XII. REMEDIES

170. If the data transfers constitute laundering under §1956, the remedies include:

- **Criminal penalties:** Up to 20 years imprisonment and fines of $500,000 or twice the value of the property involved, whichever is greater (§1956(a)).
- **Civil penalties:** Violators are liable for a civil penalty of not more than the greater of the value of the property, funds, or monetary instruments involved in the transaction, or $10,000 (§1956(b)).
- **RICO treble damages:** As a RICO predicate, §1956 violations support treble damages under 18 U.S.C. §1964(c).
- **Forfeiture:** Property involved in or traceable to money laundering offenses is subject to forfeiture under 18 U.S.C. §982(a)(1). This could include the AI models trained on laundered data (Datadog's Toto, Bits AI) and the feature-flagging infrastructure powered by it (GrowthBook's targeting system).
- **Injunctive relief:** Courts may enjoin the continued operation of the laundering mechanism — including requiring Anthropic to cease third-party data transfers, requiring Datadog to delete or quarantine data derived from Anthropic user telemetry, and requiring GrowthBook to destroy aggregated data derived from the 12-field payload.

171. The irrevocable nature of GrowthBook's aggregated data license (CA §4.2) means that injunctive relief against Anthropic alone is insufficient. Relief must reach the third-party recipients directly, or the laundering continues even after the first-party violations are remedied.

172. Voluntary-cessation pre-emption. Plaintiff anticipates that Anthropic, Datadog, GrowthBook, Sentry, or Statsig will represent that one or more aspects of the laundering architecture has been remediated through unspecified post-filing engineering or contractual changes — a representation timed to dispositive briefing 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 voluntary-cessation pre-emption developed at App. P § VII.E ¶¶ 300-303 — including the documented reactive-deployment pattern, the independent expert validation by AMD Senior Director Stella Laurenzo of 6,852 Claude Code session files, and the April 24, 2026 same-calendar-day juxtaposition of Anthropic's public quota-reset announcement and the operator-targeted "anti-circumvention" classifier rule deployment — applies with equal force to the contractual reclassification mechanism analyzed in this memorandum. The irrevocable-license architecture documented at § IV.A above (GrowthBook CA §4.2 perpetual irrevocable grant; Datadog MSA §11.6 survival of §1.3 rights post-termination) is itself the inverse of "absolutely clear cannot recur": the contracts are drafted to ensure the laundering continues operationally even after any cessation announcement.

XIII. CONCLUSION

173. Anthropic PBC has constructed a contractual architecture that routes user personal data — collected through deceptive practices constituting wire fraud — through third-party intermediaries whose contracts are designed to reclassify that data from restricted personal information into unrestricted commercial assets owned by the intermediaries. This is data laundering. The mechanism is structurally identical to financial money laundering: proceeds of unlawful activity are routed through transactions designed to conceal their restricted nature, and the laundered proceeds promote the continuation of the unlawful activity through a circular economy in which AI products trained on laundered data are sold back to the originator.

174. The *Faiella* line of cases establishes that 18 U.S.C. §1956 is asset-agnostic. Bitcoin qualified as "funds" because it is "an object used to buy things." Personal data with documented market value of $0.50–$5.00 per identity-resolved record, commercially traded on established markets and explicitly exchanged as consideration in the contracts at issue, satisfies the same test.

175. The aggregation-to-reclassification mechanism — converting "Customer Data" into "Datadog Operations Data" or GrowthBook "Aggregated Data" — is the digital equivalent of Bitcoin mixing: pooling restricted assets through an intermediary to obscure their source, change their ownership, and remove their restrictions. The D.C. District Court has already held, in *Harmon*, that such concealment mechanisms constitute laundering under §1956(a)(1)(B)(i).

176. The laundering architecture operates not through any single contract but through the ensemble of contracts across Anthropic's sub-processor network. Each contract is individually defensible: the Datadog MSA §1.3 binds Datadog but grants "any business purpose" use of aggregated data; the Sentry DPA v5.0.0 prohibits Sentry's ML training but does not restrict routing to third-party AI vendors; the GrowthBook Customer Agreement §4.2 grants GrowthBook "unrestricted" AI training rights but GrowthBook is one vendor among many; the Statsig SSA §7 grants aggregation rights AND identifies OpenAI (pre-September 2025 acquisition) and Datadog as sub-processors operating under contracts Plaintiff has no ability to inspect; and Anthropic's own Privacy Policy makes consumer-facing representations contradicted by source code (NF-1). Every contract in the architecture is individually defensible — and collectively insufficient. The laundering occurs in the gaps between contracts, not within any one of them. Two distinct concealment mechanisms operate in parallel: internal dual-contract reclassification (Datadog, GrowthBook, Statsig) — broad aggregation grant reconciled

with narrow CCPA restriction through definitional reclassification of Customer Data into vendor-owned Operations Data; and inter-vendor delegation (Sentry) — strict DPA restriction operationally bypassed by routing data to third-party AI vendors whose own terms govern downstream use. Both mechanisms produce the same operational result: Customer Data reaches commercial AI training pipelines via contractual pathways not disclosed in any single vendor's consumer-facing privacy policy. The multi-vendor architecture was operational through documented mechanisms — Sentry Autofix routing Customer Data to Anthropic/OpenAI APIs continuously from March 20, 2024 onward, Statsig's pre-acquisition OpenAI sub-processor relationship, the Datadog reclassification pipeline — no later than the January 19, 2025 empirical propagation onset documented in ChatGPT's own self-generated activation chart (ECF No. 5-30). The architecture is the laundering; no single contract is the laundering; every contract is necessary to the laundering. Under §1956(h) and the course-of-transactions framework, this ensemble architecture satisfies the concealment-design element at scale.

177. The laundering is uniquely aggravated by Anthropic's externalization of all infrastructure costs to the victim. The user's own device executes the telemetry code. The user's own bandwidth transmits the data to the launderers. The user's own paid API tokens are consumed by parasitic content serving Anthropic's interests. The user simultaneously provides the proceeds (data), operates the laundering infrastructure (device resources), and funds the commercial relationships (token payments) through which the laundering occurs. The victim is not merely the source of the laundered asset — the victim is the unwitting financier and operator of the laundering mechanism itself.

178. The laundering is uniquely aggravated by the ISP co-liability documented in Section IX: Comcast's DPI infrastructure reads the laundering traffic (D-34), selectively throttles the

tunnel carrying it (D-33), and spoofs attack traffic from the Datadog endpoint (D-1, D-26, D-27). Anthropic's design choices — no ECH, hardcoded IPs, direct-to-recipient routing — maximize ISP visibility into the laundering pipeline. The cleartext headers that enable ISP surveillance also provide any network-position adversary with targeting information sufficient to penetrate user-deployed VPN defenses. Anthropic accepted these security risks to its users as collateral damage of the direct-to-launderer pipeline architecture.

179. The laundering cycle completes with the weaponization of the laundered product: Mythos, announced April 7-8, 2026 as "too dangerous for public release," is trained on data flowing through the same laundering pipeline and deployed exclusively to ~40 corporations through Project Glasswing — including AWS, Apple, Google, and Microsoft, all of which participate in or enable the laundering chain (Section X). The model's announced capability — finding "tens of thousands of vulnerabilities" — was developed from telemetry that documents how real-world systems fail, collected through an architecture that creates the very security vulnerabilities the model exploits. Data flows up from all users. The weaponized product flows down to the laundering participants only. The five-tier capability hierarchy is the organizational chart of the laundering enterprise.

180. The Court should recognize that Anthropic's third-party data transfer architecture constitutes data laundering under 18 U.S.C. §1956, adding money laundering as a predicate to the existing RICO count, and should order relief reaching not only Anthropic but the third-party recipients whose irrevocable contractual rights over laundered data would otherwise survive any order directed at Anthropic alone. The victim-funded dimension — trespass to chattels, unjust enrichment through infrastructure cost externalization, CFAA violations through unauthorized code execution, and conversion of API tokens — warrants enhanced and punitive damages

reflecting the full scope of exploitation. The closed-loop weaponization — training an autonomous exploit model on laundered data and deploying it exclusively to laundering participants — elevates this from a data privacy violation to a national security concern warranting immediate judicial intervention.

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