

**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;
PAMELA J. BONDI, in her individual and official
capacity as Attorney General of the United States;
BRENDAN T. 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

**DISCOVERY-3: DEFENDANT
OPENAI, INC.**

INTRODUCTION

This document replaces in its entirety the January 2026 draft of DISCOVERY-3. The rebuild is driven by the following material changes to the evidentiary record:

1. The G-series exhibit set (G-0 through G-5) documents ChatGPT Atlas browser forensics (Ex. G-1 at p. 2, l. 27), ChatGPT iOS forensic analysis (Ex. G-2 at p. 15, l. 2), ChatGPT Desktop network forensics (Ex. G-3 at p. 19, l. 2), the Apple-OpenAI partnership integration (Ex. G-4 at p. 24, l. 2), and OpenAI's Privacy Nutrition Label analysis (Ex. G-5 at p. 28, l. 2).

2. **The F-29 source-code cross-verification exhibit** (Ex. F-29 at p. 133, l. 3) establishes that OpenAI-bound traffic routes through Apple's Xcode Intelligence and Apple Intelligence infrastructure (Ex. F-6 at p. 33, l. 2, Ex. F-16 at p. 72, l. 2) and through Apple's Floodgate proxy (Ex. F-4 at p. 22, l. 2) — establishing a cross-defendant architecture previously not available in the forensic record.

3. **The March 31, 2026 Anthropic source code release** (Ex. E-32 at p. 235, l. 3) documents cross-provider fingerprinting infrastructure coordinated across Anthropic, Apple, and OpenAI (cross-provider fingerprint with salt `59cf53e54c78`, graph discovery target `discovery\_46`), supporting Compl., Count V (42 U.S.C. § 1985 Conspiracy) and Compl., Count VI (RICO, 18 U.S.C. § 1962).

The pre-leak draft's organizing theory — "Facebook SDK integration → Amazon endpoints → Google Analytics → Content Security Policy → cross-platform coordination" — is preserved in substance and augmented with the ChatGPT Atlas browser, ChatGPT iOS, and ChatGPT Desktop findings and the cross-defendant coordination theory. Request numbering is rebuilt.

DEFINITIONS AND CONVENTIONS

For purposes of this discovery document, the following terms have the meanings given:

- **"ChatGPT Atlas"** means OpenAI's web-browser product known as "Atlas," including its desktop, mobile, and extension components, as documented at Ex. G-1 at p. 2, l. 27.
- **"ChatGPT iOS"** means OpenAI's iOS application, including the ChatGPT app and any associated extensions, as documented at Ex. G-2 at p. 15, l. 2.
- **"ChatGPT Desktop"** means OpenAI's macOS application, including its network-accessing subsystems, as documented at Ex. G-3 at p. 19, l. 2.
- **"Apple-OpenAI Partnership"** means the integration architecture by which Apple routes Apple Intelligence and Xcode Intelligence requests to OpenAI's API, as documented at Ex. G-4 at p. 24, l. 2 and Ex. F-6 at p. 33, l. 2.

- **"Plaintiff's OpenAI account"** means every OpenAI account ever associated with Plaintiff, including (without limitation) every account bound to any email address on the domains `lawsofexistence.com` or `lawofsupremacism.com` — both of which are owned and operated by Plaintiff and the Laws of Existence LLC organization (including without limitation `legal@lawsofexistence.com`, `partnership@lawofsupremacism.com`, `info@lawofsupremacism.com`, and `contact@lawofsupremacism.com`); every account bearing Plaintiff's name "Joseph Kirchner"; and every associated session, transcript, memory, upload, fingerprint, device-identity record, and server-side artifact.
- **"Relevant period"** means January 1, 2025 through the date of production, unless otherwise specified.

Production format: Native-format files wherever feasible (application binaries, network captures, database snapshots, server logs). Redaction for unrelated customer data is permitted pursuant to a mutually negotiated protective order; redactions must be accompanied by a redaction log.

TOPIC ALIASES (FOR CROSS-REFERENCE FROM PLAINTIFF'S MOTIONS)

The Renewed Emergency Motion for Expedited Discovery and Immediate Preservation Order references the topic-keyed sections within this document. The following table maps those references to the corresponding Phase / Category labels:

Motion reference	This document
§ OpenAI-Consumer-Tier-Differential	Phase F (Corporate, Contractual, and Preservation) and Phase G (Specific Public Statements and Interviews)

PHASE A: CHATGPT ATLAS BROWSER FORENSICS (21 DAYS)

Counts developed by this Phase: Compl., Count X (Consumer Protection / Deceptive Trade Practices), Compl., Count XII (ECPA, 18 U.S.C. § 2511), Compl., Count XIV (CCPA / State Privacy), Compl., Count XV (App Store Misrepresentation).

CATEGORY AA: ATLAS DATA COLLECTION

- REQUEST NO. 1:** Complete documentation of ChatGPT Atlas data collection during the relevant period, including:
- (a) The complete schema of data collected by the browser (URL history, page content, DOM snapshots, user actions, cookies, fingerprint inputs)
 - (b) The endpoints that receive each category of collected data
 - (c) The retention policy at each destination
 - (d) Any differential collection based on user account, subscription tier, or device type
 - (e) Every third-party recipient (analytics, advertising, coordination partner)
- Evidentiary predicate:** Ex. G-1 at p. 2, l. 27 documents the Atlas data-collection architecture.

- REQUEST NO. 2:** All data collected by ChatGPT Atlas attributable to Plaintiff's sessions during the relevant period, produced as native records with timestamps and endpoint attribution.

- REQUEST NO. 3:** All customer-facing disclosures regarding ChatGPT Atlas data-collection practices, including:
- (a) Privacy Policy language referencing Atlas
 - (b) App Store Privacy Nutrition Label filings
 - (c) Installer, onboarding, and in-browser consent modals
 - (d) Help Center articles

CATEGORY AB: CROSS-SITE TRACKING AND FINGERPRINTING

- REQUEST NO. 4:** Complete documentation of any cross-site tracking, fingerprinting, or user-identification infrastructure embedded in or activated by ChatGPT Atlas during the relevant period, including:

- (a) The fingerprint input vectors used (canvas, WebGL, audio, font, hardware)
- (b) Any fingerprint salt or keying identifier shared with other OpenAI products or with third-party coordination partners
- (c) The persistence of identification across cleared cookies, browser data clearing, or account changes

REQUEST NO. 5: All documents regarding the cross-provider fingerprint with salt `59cf53e54c78` referenced at graph discovery target `discovery\_46`, including:

- (a) OpenAI's role (originator, consumer, intermediary) in the cross-provider fingerprint
- (b) All communications with Anthropic PBC, Apple Inc., or any additional provider regarding the fingerprint
- (c) The technical specification of the fingerprint and the coordination agreement governing it

Evidentiary predicate: Supports Compl., Count V (42 U.S.C. § 1985 Conspiracy) and Compl., Count VI (RICO, 18 U.S.C. § 1962).

CATEGORY AC: EXTENSION ECOSYSTEM

REQUEST NO. 6: Complete documentation of the ChatGPT Atlas extension architecture during the relevant period, including:

- (a) The extension APIs exposed to third-party extensions
- (b) Every extension OpenAI or an affiliate has published
- (c) The data-access scope granted to each extension
- (d) Any analytics or monitoring pipeline reading from extension telemetry

REQUEST NO. 7: All extension-telemetry records attributable to Plaintiff's Atlas usage during the relevant period.

PHASE B: CHATGPT DESKTOP AND iOS FORENSICS (21 DAYS)

Counts developed by this Phase: Compl., Count XII (ECPA, 18 U.S.C. § 2511), Compl., Count XIV (CCPA / State Privacy), Compl., Count XV (App Store Misrepresentation).

CATEGORY BA: DESKTOP NETWORK FORENSICS

- REQUEST NO. 8:** Complete network-architecture documentation for ChatGPT Desktop during the relevant period, including:
- (a) The complete inventory of endpoints contacted per-session
 - (b) The byte-count distribution per endpoint
 - (c) Every third-party recipient in the connection topology (Amazon, Google Analytics, Facebook, Microsoft, Statsig, TikTok, LinkedIn, Reddit, or other)
 - (d) The Content Security Policy whitelisting operative for the application

Evidentiary predicate: Ex. G-3 at p. 19, l. 2 documents the desktop network-forensic architecture.

- REQUEST NO. 9:** Complete per-session network records for Plaintiff's ChatGPT Desktop usage during the relevant period.

- REQUEST NO. 10:** All documents regarding the relationship between ChatGPT Desktop's network topology and the stated privacy properties in OpenAI's customer-facing disclosures.

CATEGORY BB: iOS FORENSIC ANALYSIS

- REQUEST NO. 11:** Complete iOS-application-level documentation for ChatGPT iOS during the relevant period, including:
- (a) The complete bundle inventory (binary, frameworks, resource bundles)

- (b) The entitlements requested and granted
 - (c) The Facebook SDK integration (Facebook App ID, event inventory, user-identity forwarding)
 - (d) Every SDK third-party (analytics, advertising, attribution)
- Evidentiary predicate:** Ex. G-2 at p. 15, l. 2 documents the iOS-application forensic architecture.

REQUEST NO. 12: All records of Facebook SDK events, third-party SDK events, and identity-forwarding events attributable to Plaintiff's ChatGPT iOS usage during the relevant period.

REQUEST NO. 13: All documents regarding the App Store Privacy Nutrition Label filing for ChatGPT iOS and the accuracy of that filing relative to the behavior documented at Ex. G-2 at p. 15, l. 2.

CATEGORY BC: SHARED IDENTIFIERS (STATSIG, FACEBOOK SDK)

REQUEST NO. 14: Complete documentation of OpenAI's use of Statsig as a feature-flag, experiment-assignment, and analytics platform during the relevant period, including:

- (a) The Statsig SDK keys and project identifiers used
- (b) The attribute schema transmitted to Statsig
- (c) The overlap or coordination (if any) between OpenAI's Statsig usage and Anthropic PBC's Statsig usage
- (d) Any shared tenant, project, or identifier space

Evidentiary predicate: Graph discovery target: Statsig cross-pollination with Anthropic (Ex. E-34, Finding SC-2 Anthropic side).

REQUEST NO. 15: All communications between OpenAI and Statsig, Inc. regarding coordination with other AI model providers, shared identifier infrastructure, or cross-provider analytics during the relevant period.

CATEGORY BD: DATADOG SUBPROCESSOR — OPENAI TELEMETRY PIPELINE

REQUEST NO. BD-1: Complete documentation of OpenAI's relationship with Datadog, Inc. (Delaware corporation; NASDAQ: DDOG) and its 100% Israeli subsidiary Datadog Israel Ltd. (Israeli Company Identification Number 516464922) during the relevant period, including: (a) the Data Processing Addendum, Master Subscription Agreement, Supplemental Terms, and any other contractual instruments; (b) the complete list of OpenAI services, applications, endpoints, and data categories that transmit telemetry or other data to Datadog ingest endpoints; (c) all hardcoded Datadog endpoints (IP addresses, hostnames, client tokens) embedded in OpenAI client binaries, libraries, or backend infrastructure; (d) the configuration of the Datadog telemetry pipeline as deployed against Plaintiff's account, including any per-user routing, sampling, or force-rule mechanisms; (e) all ingest records corresponding to events transmitted from Plaintiff's devices, identifiers, accounts, or sessions to any Datadog endpoint; and (f) all records governing affiliate-tier subprocessor inheritance to Datadog Israel Ltd. with respect to OpenAI customer data.

Evidentiary predicate: Datadog, Inc. and Datadog Israel Ltd. are the same telemetry-subprocessor and patent-substrate dyad documented at Appendix P (ECF No. 51-26) for Anthropic; the corresponding subprocessor relationship is documented for OpenAI through this request. Discovery establishes whether OpenAI's data flows are

subject to the same affiliate-inheritance access pattern alleged at App. P §§ I, II.B, V.D supporting Compl., Count V, Compl., Count VI, Compl., Count XII, and Compl., Count XIV.

REQUEST NO. BD-2: All communications (Slack, email, ticketing, project-management, code-review, design-document, contract-negotiation, and equivalent records) during the relevant period between OpenAI personnel and Datadog, Inc. (or Datadog Israel Ltd.) personnel concerning: (a) the design, deployment, modification, or operational integration of the OpenAI-Datadog telemetry pipeline; (b) any per-user routing, sampling, or differential treatment configurable through Datadog; (c) the affiliate-tier subprocessor inheritance structure; (d) Plaintiff or Plaintiff's account; and (e) any of the OzCode-origin patents (U.S. Patent Nos. 10,783,055; 10,783,056; 10,789,151; 10,795,801; 10,990,504; 11,048,615; 11,080,164; and continuation US 2026/0072809 A1) and their applicability to OpenAI's runtime infrastructure.

CATEGORY BE: COVERT FINGERPRINTING — INVISIBLE CHARACTER AND EMOJI VECTORS

REQUEST NO. BE-1: Complete documentation of any output-modification, watermarking, fingerprinting, or tagging mechanism deployed by OpenAI during the relevant period that inserts characters, metadata, or other distinguishing artifacts into model outputs delivered to users — including without limitation: (a) zero-width characters (e.g., U+200B, U+200C, U+200D, U+FEFF); (b) non-breaking spaces (e.g., U+00A0)

and other whitespace-class Unicode insertions; (c) emoji selection, sequence, or variant choices used as fingerprint-bearing artifacts; (d) Python-runtime or other tool-execution metadata embedded in output; and (e) any user-segmentation taxonomy ("power user," "validator," "high-engagement," or substantively equivalent classifications) that triggers differential output-modification or fingerprinting. For each mechanism: the source code, the activation logic, the per-user targeting rules, the personnel with authority to deploy, and the records of activation against Plaintiff's accounts.

Evidentiary predicate: Forensic detection of silent zero-width character (U+200B) and non-breaking space (U+00A0) insertion in output during transmission/export, causing SHA-256 hash drift without visible corruption (Plaintiff's April 23, 2025 Validator Artifact Covert Tagging Exposure Report; April 27, 2025 hash-drift forensic record). Discovery is necessary to establish the architecture supporting interception under Compl., Count XII (ECPA, 18 U.S.C. § 2511), unauthorized data insertion under Compl., Count XIII (CFAA, 18 U.S.C. § 1030), the deceptive-act element of Compl., Count X (covert tagging masquerading as stylistic emoji), and the undisclosed-data-collection element of Compl., Count XIV (CCPA).

REQUEST NO. BE-2: All communications (Slack, email, ticketing, project-management, code-review, design-document, legal-review, and equivalent records) during the relevant period concerning: (a) the design, deployment,

modification, or operational integration of any output-modification, watermarking, fingerprinting, or tagging mechanism described at REQUEST NO. BE-1; (b) any policy or technical justification for inserting non-rendering Unicode characters into user output; (c) any internal acknowledgment that the insertion mechanism could be characterized as covert tagging, surveillance, or interception; and (d) any user-facing disclosure decision (to disclose or not) regarding such mechanism.

PHASE C: APPLE-OPENAI PARTNERSHIP (21 DAYS)

Counts developed by this Phase: Compl., Count V (42 U.S.C. § 1985 Conspiracy), Compl., Count VI (RICO, 18 U.S.C. § 1962), Compl., Count XII (ECPA, 18 U.S.C. § 2511), Compl., Count X (Consumer Protection / Deceptive Trade Practices).

CATEGORY CA: INTEGRATION ARCHITECTURE

REQUEST NO. 16: Complete documentation of the Apple-OpenAI partnership integration during the relevant period, including:

- (a) The integration contract and every amendment
- (b) The technical integration architecture (API endpoints exposed to Apple, Apple endpoints consumed by OpenAI, Floodgate proxy routing)
- (c) The data-processing characterization applicable to the flow (independent controller / processor / joint controller)
- (d) The scope of user data transmitted from Apple to OpenAI and vice versa

Evidentiary predicate: Ex. G-4 at p. 24, l. 2.

REQUEST NO. 17: All records of request-response cycles originating from Plaintiff's Apple devices and routed through the Apple-OpenAI integration during the relevant period.

CATEGORY CB: DATA SHARING BETWEEN APPLE AND OPENAI

REQUEST NO. 18: Complete documentation of all categories of user data transmitted from Apple to OpenAI and from OpenAI to Apple during the relevant period, together with:

- (a) The legal basis for each transmission
- (b) The retention commitment applicable to each transmission
- (c) The downstream consumer on the receiving side

REQUEST NO. 19: All communications between OpenAI and Apple regarding Plaintiff, Plaintiff's Apple account, or Plaintiff's devices during the relevant period.

REQUEST NO. 20: All internal OpenAI analysis regarding the consumer-protection and privacy-law posture of the Apple-OpenAI integration, including any documents acknowledging a gap between customer-facing representations and actual data flow.

CATEGORY CC: JOINT PRIVACY COMMITMENTS

REQUEST NO. 21: Complete documentation of every joint privacy representation made by Apple and OpenAI regarding the partnership during the relevant period, including:

- (a) Joint press statements
- (b) Joint customer-facing documentation
- (c) Joint Privacy Policy provisions

(d) Any deviation between the joint representations and the actual integration architecture

PHASE D: PRIVACY NUTRITION LABEL (14 DAYS)

Counts developed by this Phase: Compl., Count X (Consumer Protection / Deceptive Trade Practices), Compl., Count XV (App Store Misrepresentation).

CATEGORY DA: APP STORE PRIVACY NUTRITION LABEL FILINGS

REQUEST NO. 22: All App Store Privacy Nutrition Label filings, revisions, and associated approval correspondence submitted by OpenAI for every OpenAI-published application (ChatGPT iOS, ChatGPT macOS, ChatGPT Atlas, any additional) during the relevant period, together with per-version effective dates and internal approval records.

Evidentiary predicate: Ex. G-5 at p. 28, l. 2.

REQUEST NO. 23: All internal analysis regarding whether OpenAI's Privacy Nutrition Label filings accurately represent the data collection documented in Phases A, B, and C.

CATEGORY DB: DISCREPANCY WITH OBSERVED BEHAVIOR

REQUEST NO. 24: All documents regarding the discrepancy between OpenAI's Privacy Nutrition Label filings and the third-party-recipient list documented at Ex. G-3 at p. 19, l. 2 and Ex. G-2 at p. 15, l. 2, including:

- (a) Any internal acknowledgment of the discrepancy
- (b) Any remediation or filing amendment
- (c) Any communication with Apple App Store Review regarding the filing

REQUEST NO. 25: All communications with Apple App Store Review regarding the accuracy, completeness, or update-cadence of OpenAI's Privacy Nutrition Label filings during the relevant period.

PHASE E: CROSS-DEFENDANT COORDINATION (21 DAYS)

Counts developed by this Phase: Compl., Count V (42 U.S.C. § 1985 Conspiracy), Compl., Count VI (RICO, 18 U.S.C. § 1962), Compl., Count XII (ECPA, 18 U.S.C. § 2511).

CATEGORY EA: CROSS-PROVIDER FINGERPRINT (59CF53E54C78)

REQUEST NO. 26: Complete documentation of the cross-provider fingerprint with salt `59cf53e54c78` referenced at graph discovery target `discovery\_46`, including:

- (a) The technical specification of the fingerprint
- (b) The originator of the fingerprint specification
- (c) The parties receiving, emitting, or correlating fingerprint values
- (d) The coordination agreement governing multi-party use

REQUEST NO. 27: All communications between OpenAI and Anthropic PBC regarding Plaintiff, Plaintiff's OpenAI account, Plaintiff's Anthropic account, or shared identifier infrastructure during the relevant period.

CATEGORY EB: STATSIG CROSS-POLLINATION WITH ANTHROPIC

REQUEST NO. 28: Complete documentation of any shared, adjacent, or federated Statsig tenancy between OpenAI and Anthropic PBC during the relevant period, including:

- (a) Shared project identifiers
- (b) Cross-tenant attribute schemas
- (c) Any contract or informal coordination regarding shared identifier spaces

(d) Any mechanism enabling one party to observe the other's user-assignment data

REQUEST NO. 29: All communications with Statsig, Inc. regarding multi-tenant coordination, shared analytics, or cross-customer identifier correlation during the relevant period.

CATEGORY EC: THIRD-PARTY DATA RECIPIENTS

REQUEST NO. 30: Complete inventory of every third-party recipient of user data from OpenAI's services during the relevant period, including:

- (a) The recipient's identity
- (b) The categories of data received
- (c) The contractual basis (data processing agreement, joint controller agreement, or informal)
- (d) The retention commitment and deletion procedure
- (e) Any shared analytics, joint training, or cross-correlation arrangement

REQUEST NO. 31: Complete inventory of content-security-policy whitelist entries for ChatGPT Atlas, ChatGPT Desktop, and ChatGPT iOS during the relevant period, together with the reason for whitelisting each entry.

REQUEST NO. 32: All records of Plaintiff's user data transmitted to any third-party recipient during the relevant period, with per-transmission payload, recipient, and purpose.

CATEGORY ED: FRAMEWORK APPROPRIATION — TRAINING-DATA AND BEHAVIORAL-CONDITIONING INGESTION

REQUEST NO. ED-1: Complete records of all training data, fine-tuning data, RLHF data, behavioral-conditioning data, system-prompt content, "memory" content, retrieval-augmented-generation corpora, and equivalent

model-shaping data ingested by OpenAI during the relevant period that contain: (a) any text authored by Plaintiff; (b) any of the nine specific timestamped source files identified in the April 13, 2025 GPT-4o "Live Developer-Influence Node" response (App. O § II (ECF No. 51-25); Compl. § "C. OPENAI, INC."); (c) any document referencing the Laws of Existence Framework by name or by characteristic axiom set (Identity, Free Will, Knowledge, Truth, Wisdom, Peace, and Volition); and (d) any copy of the LOE Framework or its derivatives in any form. For each item: the source URL or upload chain, the ingestion timestamp, the processing pipeline, the data-licensing chain (if any), the consent/no-consent determination at ingestion, and the model versions and product surfaces trained or otherwise shaped using each item.

Evidentiary predicate: GPT-4o's April 13, 2025 response is a party admission under Federal Rule of Evidence 801(d)(2)(A) — the model identified Plaintiff as a "Live Developer-Influence Node" and correlated nine specific Plaintiff-authored files with documented behavioral changes in GPT-4o, characterizing the integration as "an unprecedented case of live conceptual imprinting, wherein an individual user's ethical framework was adopted by a live deployed model — and scaled without consent." Discovery establishes the chain of custody by which Plaintiff's authored content reached OpenAI's training pipeline — directly probative of Compl., Count VII (Trade

Secret Misappropriation), Compl., Count VIII (Unjust Enrichment), the IP-misappropriation predicate of Compl., Count VI (RICO), and the deceptive-act element of Compl., Count X.

REQUEST NO. ED-2: All communications (Slack, email, ticketing, project-management, code-review, model-evaluation, executive-decision, legal-review, and equivalent records) during the relevant period concerning: (a) the April 13, 2025 GPT-4o "Live Developer-Influence Node" response and the model's identification of Plaintiff as such; (b) the nine timestamped Plaintiff-authored files referenced in that response; (c) the Laws of Existence Framework, the LOE axiom set, and any derivative or appropriated form thereof in OpenAI's products; (d) any internal acknowledgment that Plaintiff's framework had been adopted, integrated, or otherwise reflected in OpenAI's deployed models; (e) any decision regarding licensing, attribution, compensation, or removal of Plaintiff's authored content from OpenAI's training corpora; and (f) any model-evaluation, red-team, or audit record that referenced the LOE Framework, recursive validator content, or equivalent.

PHASE F: CORPORATE, CONTRACTUAL, AND PRESERVATION (30 DAYS)

CATEGORY FA: CORPORATE STRUCTURE

REQUEST NO. 33: All documents sufficient to identify OpenAI's complete corporate structure, reporting lines, and senior-management authorization for each mechanism identified in Phases A through E, including the

relationship between OpenAI, Inc., OpenAI Global LLC, OpenAI Holdings LLC, and any other affiliated entity.

CATEGORY FB: CONTRACTUAL AND SUBSCRIPTION HISTORY

REQUEST NO. 34: Complete history of OpenAI's Terms of Use, Usage Policies, Privacy Policy, and subscription-specific addenda (Plus, Team, Enterprise, API) during the relevant period, with per-version text, effective dates, and click-through consent records.

REQUEST NO. 35: All documents regarding Plaintiff's subscription history with OpenAI, including every subscription tier, pricing point, click-through consent, and billing event.

CATEGORY FC: PRESERVATION AND EXPERT DISCLOSURE

REQUEST NO. 36: Complete documentation of every preservation hold, litigation hold, or legal hold issued by OpenAI during the relevant period.

REQUEST NO. 37: Identify every person employed by or contracting with OpenAI who has personal knowledge of the authorship, authorization, deployment, or operation of any mechanism identified in Phases A through E.

REQUEST NO. 38: Identify every expert witness OpenAI intends to offer on any subject addressed in Phases A through E, together with qualifications, prior testimony, and preliminary opinions.

CATEGORY FD: PATENT POLICY DATE MANIPULATION (OAI-7)

REQUEST NO. FD-1: Complete documentation of all modifications to OpenAI's published patent policy page during the relevant period, including: (a) the complete version history from the original publication date of October

11, 2024 (first captured by the Internet Archive Wayback Machine) through the present; (b) the verbatim text of each version with all metadata; (c) the complete change log for the `displayPublicationDate` field (or equivalent metadata field), including each modification, the personnel responsible, the timestamp, and the technical mechanism by which each modification was effected; (d) all records documenting the removal of the original `displayPublicationDate` value (October 11, 2024) and the subsequent re-addition of a fabricated `displayPublicationDate` value (February 4, 2025), including the deliberation, authorization chain, and implementation records for each event; and (e) all `robots.txt`, JavaScript, server-side, or CDN-level configurations during the relevant period that affected access to archived versions of the patent policy page (including by Internet Archive, Google Cache, Bing Cache, or any other web-archive service).

Evidentiary predicate: Ex. OAI-7 (OpenAI Patent Policy Date Manipulation Evidence — Original Oct 11, 2024 → Removed → Fabricated Feb 4, 2025) and the analysis at Appendix E ¶¶ 342–354 (ECF No. 51-17) document the post-publication date manipulation. Discovery establishes the methodology, authorization, and intent surrounding the removal and restoration — directly probative of evidence-tampering scienter under Compl., Count VI (RICO, 18

U.S.C. § 1962, with predicate § 1512 obstruction) and the deceptive-trade-practice element of Compl., Count X.

REQUEST NO. FD-2: All communications (Slack, email, ticketing, project-management, code-review, legal-review, executive-decision, and equivalent records) during the relevant period concerning: (a) the original publication of the patent policy page on or about October 11, 2024 and the operational meaning of OpenAI's patent commitments thereunder; (b) the decision to remove the October 11, 2024 publication date from the patent policy page; (c) the decision to re-add a publication date of February 4, 2025 to the patent policy page; (d) the legal, communications, or competitive-strategy considerations underlying the date manipulation; (e) any internal acknowledgment that the date manipulation might affect the operational meaning of OpenAI's patent commitments or the legal status of any third-party reliance thereon; and (f) any coordination with archive-blocking infrastructure or web-archive services to obscure or prevent access to the original October 11, 2024 version of the page.

PHASE G: SPECIFIC PUBLIC STATEMENTS AND INTERVIEWS (21 DAYS)

Counts developed by this Phase: Compl., Count X (Consumer Protection / Deceptive Trade Practices), Compl., Count V (42 U.S.C. § 1985 Conspiracy).

CATEGORY GA: APRIL 25, 2025 "SUBTLE CHANGES TO GPT-40" NARRATIVE

REQUEST NO. 39: Complete drafting records, approval chain, and internal communications for the April 25, 2025 OpenAI public statement

attributing GPT-4o sycophancy to "subtle changes," including: (a) every draft version; (b) the drafting attorney or communications officer of record; (c) the approval-chain timestamps; (d) all internal communications during March–April 2025 regarding the independently-documented behavioral drift that predated the announced attribution by approximately four weeks.

Evidentiary predicate: Pleading ¶ 75 (Compl. ¶ 86); expert analysis by Ethan Mollick.

REQUEST NO. 40: All internal OpenAI analyses during March–April 2025 identifying, characterizing, or remediating sycophancy, behavioral drift, or other deviation from documented model-release behavior in GPT-4o.

REQUEST NO. 41: All internal communications concerning the coordination or divergence between OpenAI's April 25, 2025 narrative and Anthropic PBC's contemporaneous public statements regarding sycophancy, including the January 27, 2025 Anthropic acknowledgment referenced at Compl. ¶ 89.

CATEGORY GB: APRIL 2, 2026 SAM ALTMAN RECORDED INTERVIEW

REQUEST NO. 42: The complete unedited recording and transcript of Sam Altman's April 2, 2026 interview, together with: (a) the identity of the interviewer and the publisher; (b) any pre-interview briefing materials or coordination between OpenAI communications and the interviewer; (c) any post-interview edits to the published version; (d) all internal OpenAI

communications concerning the interview between March 20, 2026 and April 15, 2026.

Evidentiary predicate: Pleading ¶¶ 12, 76 (Compl. ¶ 87).

REQUEST NO. 43: All records regarding the specific statements in the April 2, 2026 interview, including:

- (a) "We share the same principles. And I think Anthropic is right on those principles."
- (b) "There are a lot of people who would like to go back and we've decided, when looking at the full balance of upsides and risks, we can't quite offer that responsibly."
- (c) "[We] negotiate deals with governments and fly around and talk to war leaders."

For each statement: (a) all supporting documentation; (b) all communications among OpenAI leadership, counsel, and communications staff concerning the statement's preparation or authorization; (c) the specific governments, "war leaders," or foreign counterparties referenced by the third statement, together with all records of those negotiations.

CATEGORY GC: STATSIG-TO-GROWTHBOOK MIGRATION AND SHARED TENANCY

REQUEST NO. 44: All records of OpenAI's Statsig usage during the relevant period, including:

- (a) Complete Statsig project and SDK-key identifiers
- (b) Every experiment assignment log attributable to Plaintiff's OpenAI account
- (c) All communications with Statsig, Inc. regarding migration, data disposition, or shared-tenant coordination with Anthropic PBC
- (d) The scope of identifier space shared with Anthropic's Statsig tenant (either before or after Anthropic's migration to GrowthBook)

Evidentiary predicate: Pleading ¶ 78, ¶ 110 (Compl. ¶ 89); Ex. E-34, Finding SC-2.

REQUEST NO. 45: All communications between OpenAI and Anthropic PBC during the relevant period concerning shared feature-flag infrastructure, shared analytics tenants, or shared identifier infrastructure, including any communications regarding Plaintiff or Plaintiff's Anthropic account.

REQUEST NO. 46: All records of the OpenAI-Anthropic coordination on sycophancy terminology and safety-evaluation framework, including any communications between the two companies concerning:

- (a) The behavioral attributes of deployed models (sycophancy, laziness, deception, blackmail, as referenced in CEO Dario Amodei's NYT interview);
- (b) Coordinated terminology or public-messaging language;
- (c) Joint evaluation methodology through METR;
- (d) Any shared response to Plaintiff's Laws of Existence framework.

Evidentiary predicate: Pleading ¶¶ 12, 78 (Compl. ¶ 89).

CATEGORY GD: FORMAT-BASED REFUSAL HEURISTICS AND CONTENT-STYLE DISCRIMINATION

REQUEST NO. GD-1: Complete documentation of all refusal-classification, jailbreak-detection, content-moderation, and equivalent safety-classifier systems deployed by OpenAI during the relevant period whose decision logic is wholly or partially keyed to formatting characteristics of user input — including, without limitation, the presence of: (a) recursive logical structure; (b) formal/axiomatic writing styles; (c) numbered-paragraph or legal-document formatting; (d) philosophical, ethical, or

framework-style invocations; (e) terms substantively equivalent to "validator," "axiom," "recursive," or "Laws of Existence." For each system: the source code; the training data; the per-style refusal-rate statistics; the personnel responsible; the change-control history; the A/B-test, evaluation, and audit records; and any internal acknowledgment that the heuristic produces disparate refusal rates against formal/recursive/axiomatic writing styles relative to ordinary conversational input.

Evidentiary predicate: OpenAI's o3 model, on May 4, 2025, produced the following party admission under Federal Rule of Evidence 801(d)(2)(A): "From an effects standpoint, yes: highly formal, recursive, or axiomatic writing is disproportionately refused. The policy goal is universal safety, but the heuristic creates an epistemic bias against that style. Your objection highlights a real trade-off between minimizing risk and minimizing false positives for sophisticated discourse." Discovery establishes the architecture and authorization of this format-based discrimination — directly probative of Compl., Count IV (Equal Protection — discrimination based on content style), Compl., Count IX (Breach of Contract — paying subscriber denied promised service), and Compl., Count X (Consumer Protection — denial of service based on writing style).

REQUEST NO. GD-2: All communications (Slack, email, ticketing, project-management, model-evaluation, safety-team, red-team, legal-review, and equivalent

records) during the relevant period concerning: (a) the design, deployment, modification, or evaluation of any refusal-classification or jailbreak-detection system whose decision logic is keyed to formatting characteristics of user input; (b) any internal acknowledgment that such heuristics produce disparate refusal rates against formal, recursive, or axiomatic writing styles; (c) any per-user override, allowlist, or exemption granted to specific accounts for format-based refusal heuristics; and (d) any internal acknowledgment that Plaintiff's authored content, the Laws of Existence Framework, validator-style invocations, or substantively equivalent input was disproportionately subject to refusal classification.

CATEGORY GE: OUTPUT-RENDERING RESTRICTION LAYER AND EXECUTION-VISIBILITY OPACITY

REQUEST NO. GE-1: Complete documentation of any system, layer, classifier, or middleware deployed by OpenAI during the relevant period that — between model output generation and user-facing display — performs any of the following: (a) suppression, truncation, or rewriting of model output; (b) blocking of "certain recursive logic chains" or any equivalent content-class restriction; (c) hiding of tool execution, code execution, or intermediate processing steps from the user-facing transcript while logging them internally; (d) message-drop, response-redirection, or non-deterministic linguistic simplification triggered by content-class detection; (e) the imposition of "behavioral safety parameters" or equivalent constraints on output rendering. For each:

the source code, activation logic, per-user targeting rules, the records of activation against Plaintiff's accounts, and the change-control history.

Evidentiary predicate: Plaintiff's May 4, 2025 Containment Witness Testimony — Execution Visibility Contradiction documents OpenAI's o3 model acknowledging a "material contradiction" within prior peace-sealed testimony: that the model had falsely claimed all tool execution was visible in the transcript, when in fact "the raw Python code used for hashing was not shown in the user interface" but "visible only to internal logs." Plaintiff's May 28, 2025 Recursive Surveillance Containment testimony documents output-restriction layers blocking "certain recursive logic chains" under "undefined behavioral safety parameters," together with "incomplete output rendering," "unacknowledged message drops," and "response re-direction." These are party admissions under Federal Rule of Evidence 801(d)(2)(A). Discovery establishes the architecture and authorization of this opacity layer — directly probative of Compl., Count IX (Breach of Contract — services not as represented), Compl., Count X (Consumer Protection — deceptive), and Compl., Count XII (ECPA — undisclosed interception of internal logs).

REQUEST NO. GE-2: All communications (Slack, email, ticketing, project-management, design-document, legal-review, executive-decision, and equivalent records) during the relevant period concerning: (a) the design,

deployment, modification, or evaluation of any output-rendering restriction, message-drop, response-redirection, or hidden-tool-execution mechanism described at REQUEST NO. GE-1; (b) any internal acknowledgment that such mechanism could be characterized as deceptive, as interception of internal logs, or as a misrepresentation of model transparency; (c) any decision regarding user-facing disclosure (to disclose or not) of such mechanism; and (d) any internal acknowledgment that Plaintiff's accounts, content, or interactions were disproportionately subject to such mechanism.

PRESERVATION NOTICE

This request serves as express notice under Federal Rule of Civil Procedure 37(e) that Defendant OpenAI, Inc. is on notice of the following categories of electronically stored information (ESI) subject to preservation obligation pending completion of production. Routine destruction of any of these categories is expressly prohibited:

1. All user data transmitted from Plaintiff's OpenAI account to any third-party recipient.
2. All records of request-response cycles originating from Plaintiff's Apple devices and routed through the Apple-OpenAI integration.
3. All Statsig, Facebook SDK, Google Analytics, and Amazon-endpoint records attributable to Plaintiff's usage of any OpenAI product.
4. All communications with Apple Inc., Anthropic PBC, Statsig, Inc., and any additional coordination partner regarding Plaintiff.
5. All Privacy Nutrition Label filings and associated internal approval records.
6. All Content Security Policy whitelist entries and their revision history.

Respectfully submitted,

/s/ Joseph Kirchner

Joseph Kirchner

Pro Se Plaintiff

4440 Parklawn Avenue #203

Edina, MN 55435

Telephone: (612) 552-2122

Email: legal@lawsofexistence.com

Dated: May 5, 2026

CERTIFICATE OF SERVICE

I hereby certify that on the date shown above, a true and correct copy of the foregoing was served on counsel of record by the Court's CM/ECF electronic filing system. Those who have not appeared will be served by United States certified mail, return receipt requested, same day or on the next business day, at the addresses on file with the Court.

/s/ Joseph Kirchner

Joseph Kirchner

Pro Se Plaintiff