Investigative Audit Report on Frontier Model Governance, Corporate Restructuring, and Intellectual Property Provenance

Part I: The Commercial Structuring and Partner Exfiltration Paradigm

Corporate Shell Evolution and Charter Arbitrage

The corporate architecture of OpenAI represents an unprecedented restructuring of a charitable enterprise into a heavily capitalized, commercially driven entity, establishing a sophisticated framework for regulatory arbitrage and the commercialization of synthetic training data. Originally founded in 2015 as a 501(c)(3) non-profit organization with the explicit mission to ensure that artificial general intelligence benefits all of humanity, the organization fundamentally altered its trajectory in 2019 by creating a capped-profit subsidiary, OpenAI Global, LLC[1]. This holding structure was ostensibly designed to attract the vast capital reserves required for frontier model compute, while remaining under the absolute governance of the non-profit board via a manager entity, OpenAI GP LLC[2].

However, throughout 2024 and 2025, OpenAI initiated a sweeping recapitalization to transition the for-profit arm into a Public Benefit Corporation (PBC), OpenAI Group PBC, effectively dissolving the previous profit caps that had limited early investor returns to one hundred times their initial investment[3]. The mechanics of this transition bypass the original non-profit fiduciary mandates by transferring the locus of operational control, intellectual property, and enterprise data licensing entirely to the PBC. While the non-profit OpenAI Foundation retains a 26 percent equity stake—valued at approximately $130 billion based on a $150 billion total enterprise valuation—and holds special voting rights to appoint the board of directors, critics argue this structure is inherently conflicted[1].

The California Attorney General's office subjected this restructuring to intense review to determine whether the conversion constituted a failure to protect charitable assets, drawing parallels to the historic conversion of Blue Cross California[3]. Coalitions such as EyesOnOpenAI and Public Citizen have formally petitioned the Attorney General, arguing that the non-profit assets, primarily consisting of operational trade secrets, intellectual property, and the foundational algorithms themselves, were functionally exfiltrated into the for-profit vehicle without adequate disgorgement or compensation for the eliminated profit caps and control premiums[3]. Furthermore, the restructuring solidifies deep integrations with primary compute partners, notably Microsoft, enabling the seamless licensing of proprietary enterprise data through opaque data pipelines while insulating the executive leadership under the guise of the remaining non-profit charter.

Whistleblower Containment and Non-Disparagement Mechanisms

To maintain strict control over these data pipelines, dataset origin irregularities, and safety threshold overrides, OpenAI historically deployed highly aggressive whistleblower containment strategies that directly conflict with federal securities regulations. In July 2024, anonymous whistleblowers filed a formal complaint with the United States Securities and Exchange Commission, alleging that OpenAI's offboarding agreements blatantly violated SEC Rule 21F-17(a) and the Dodd-Frank Act[7]. Rule 21F-17(a) strictly prohibits any individual or entity from taking action to impede someone from communicating directly with the SEC regarding potential securities law violations[8].

The investigative evidence indicates that departing personnel, particularly technical and safety researchers possessing documented evidence of model alignment failures, were compelled to sign restrictive non-disparagement agreements and non-disclosure clauses[7]. The enforcement mechanism for these agreements was exceptionally punitive: refusal to sign the restrictive exit documents triggered a severe equity forfeiture clause, threatening the cancellation of all vested stock options[7]. By forcing employees to choose between reporting regulatory non-compliance and forfeiting immense financial equity, OpenAI constructed an effective neutralization barrier against internal dissent[7].

While executive leadership subsequently issued public apologies—stating they were unaware of the provision and that vested equity had never actually been clawed back—the chilling effect on the workforce was absolute[9]. The SEC has previously established clear precedents in cases such as *In the Matter of AppFolio, Inc.* and *In the Matter of a.k.a. Brands*, heavily penalizing corporations that require departing employees to waive their rights to whistleblower incentives or that discourage communication with federal authorities[13]. The structural strategy of leveraging vested equity as a compliance mechanism ensures that critical information regarding trade secret misappropriation and safety architecture failures remains entirely contained within the corporate ecosystem.

Whistleblower Containment Mechanism	Operational Function	Regulatory Conflict
Vested Equity Forfeiture Clauses	Threatens cancellation of earned stock options if offboarding non-disparagement agreements are refused.	Directly violates SEC Rule 21F-17(a) by imposing financial penalties on potential whistleblowers[8].
Non-Disparagement & NDA Bundling	Prohibits departing personnel from discussing dataset	Conflicts with Dodd-Frank protections; previously penalized by

	irregularities, safety overrides, or the terms of the NDA itself.	the SEC in similar corporate enforcement actions[7].
Incentive Waiver Requirements	Compels employees to disclaim and waive any right to share in monetary awards resulting from federal investigations.	Explicitly prohibited under SEC whistleblower incentive protections intended to encourage the reporting of securities violations[7].

Enterprise Supply Chain Ingestion and Private Safety Processing

The commercialization of frontier models requires the continuous ingestion of proprietary corporate workflows, logistics data, and commercial client communications processed through enterprise API tiers. To assure corporate clients of data sovereignty while simultaneously maintaining necessary algorithmic safety monitoring, OpenAI utilizes a bifurcated data retention architecture. By default, the enterprise API platform retains inputs and outputs for thirty days to resolve incidents and investigate platform abuse[15]. For highly regulated sectors, the company offers a Zero Data Retention policy, explicitly guaranteeing that customer content is excluded from abuse monitoring logs and is not utilized for generative model training[17]. However, the necessity of monitoring complex, multi-session abuses—such as coordinated threat actors probing safeguards or autonomous agents drifting from user intent—renders strict Zero Data Retention highly problematic for frontier model safety[18]. To reconcile this, OpenAI introduced Private Safety Processing, a system designed to enable offline, automated safety review without retaining customer prompts in plaintext[19]. Under Private Safety Processing, customer content is stored within customer-controlled cloud infrastructure, such as Amazon Web Services or Microsoft Azure, and doubly encrypted utilizing OpenAI-managed keys and customer-managed Enterprise Key Management systems[20].

While marketed as a cryptographic safeguard, this architecture establishes an advanced, automated ingestion vector. Although human personnel are structurally barred from accessing the plaintext data, the encrypted records are processed asynchronously within hardware-attested safety runtimes that evaluate the proprietary workflows to extract bounded safety signals and operational metadata[19]. This allows the platform to systematically evaluate corporate client communications and logistics data under the guise of automated alignment checks, maintaining a competitive advantage over rivals who mandate plaintext retention for safety monitoring, while effectively integrating global enterprise supply chains into the model's telemetry ecosystem.

Part II: Ecosystem Monopolization and Federal Contract Arbitrage

Search Indexing, Zero-Attribution Data Absorption, and the Robots.txt

Illusion

The foundation models developed by Alphabet Inc. and Google DeepMind, specifically the Gemini series, rely upon the systematic, web-scale harvesting of open content, local business registries, and proprietary database records. This zero-attribution data absorption has triggered a wave of global litigation and severe regulatory backlash from international competition authorities. In the United States, a coalition of major publishers and authors initiated class-action litigation, including *In re Google Generative AI Copyright Litigation* and *Leovy v. Google*, alleging direct and willful copyright infringement arising from the scraping of millions of copyrighted works to train generative products[21].

The regulatory response in Europe has been notably punitive. In March 2024, the French Autorité de la concurrence imposed a €250 million fine on Google for breaching commitments previously established with press publishers and news agencies[23]. The investigation revealed that Google systematically ingested content from the domains of press publishers to train the foundation model of its artificial intelligence chatbot without attribution, licensing, or compensation[23]. Crucially, the Autorité found that Google failed to provide a viable technical solution allowing publishers to opt out of artificial intelligence training without simultaneously destroying their visibility in standard search indexing[23].

In response to mounting scrutiny, Google introduced a robots.txt control token designated as Google-Extended[25]. However, technical analysis reveals that this mechanism functions largely as an optical concession rather than a robust safeguard against data extraction.

Google-Extended is not an independent web crawler; it is merely a policy flag that site owners can append to their robots.txt files to request that their content be excluded from Gemini training and grounding[26]. Because compliance with robots.txt directives is entirely voluntary, and because blocking the Google-Extended token does not prevent the primary Googlebot crawler from fetching and retaining the page content for search indexing, the mechanism fails to provide absolute data sovereignty to original creators[27]. Consequently, proprietary data continues to be absorbed into the ecosystem, with enforcement relying solely on the platform's internal goodwill rather than cryptographic or legal guarantees.

Local Marketplace Integrity and Ad Arbitrage Persistence

The integration of generative artificial intelligence into local search algorithms has exacerbated existing vulnerabilities in marketplace integrity, permitting sophisticated ad arbitrage and entity-splitting tactics to persist alongside highly monetized enterprise search products. In August 2024, the Federal Trade Commission announced a final rule strictly banning fake reviews, testimonials, and the use of artificial intelligence to generate fraudulent consumer feedback[29]. The rule explicitly prohibits businesses from misrepresenting that a website they control provides independent reviews and bans the procurement of fake indicators of social media influence[29].

Despite the implementation of this federal framework, the structural mechanics of Google's Local Services Ads remain highly susceptible to exploitation. Fraudulent operators utilize virtual office addresses and hijack legitimate business profiles to generate unverified lead flow[31]. While Google has implemented harsher profile restriction policies—including the unpublishing of

existing reviews and the application of warning labels to profiles caught utilizing fake engagement—the underlying platform mechanics still permit widespread entity-splitting[31]. Scammers frequently create secondary limited liability companies to protect their primary profile ratings from automated fraud removal, ensuring that unverified lead generation continues to cross-subsidize the broader advertising ecosystem. The enforcement metrics detailing the mitigation of these review manipulations remain opaque, shifting the burden of discovery and reporting entirely onto consumers and legitimate competitors.

Federal Infrastructure and Energy Contract Subsidization

To sustain the colossal compute power required to train and operate frontier models, Alphabet has strategically leveraged federal infrastructure and energy markets, intertwining commercial artificial intelligence deployment with public-interest utility grids. In a landmark strategic maneuver, Google executed a master development agreement with Kairos Power to deploy a fleet of small modular nuclear reactors designed to supply 500 megawatts of clean, uninterrupted electricity to its artificial intelligence data centers by the year 2030[33]. These advanced molten fluoride salt-cooled reactors, utilizing TRISO fuel, will directly power data center campuses operating on the Tennessee Valley Authority grid across Tennessee and Alabama[34].

This agreement represents a profound cross-subsidization between commercial search advertising, cloud infrastructure expansion, and federal energy contracts. The United States Department of Energy explicitly supports these deployments, citing the necessity of demand aggregation from large power consumers to ensure domestic artificial intelligence dominance and energy leadership[34]. However, by tethering the infrastructural backbone of a massive federal utility to proprietary data center demands, the corporation effectively stabilizes the staggering energy costs of synthetic training against the volatility of standard commercial energy markets[36]. This federal contract arbitrage ensures that the physical infrastructure required to maintain ecosystem monopolies is heavily underwritten by public-interest energy agreements, while the resulting foundational models remain strictly proprietary.

Infrastructure & Ecosystem Vector	Operational Mechanism	Regulatory & Market Implication
Data Absorption (Google-Extended)	A robots.txt token separating AI training opt-outs from standard search indexing[25].	Creates an illusion of control; scraping continues via standard crawlers, exposing publishers to zero-attribution ingestion[27].
Local Ad Arbitrage	Entity-splitting and virtual office hijacking within	Circumvents FTC final rules on fake reviews,

	Local Services Ads[31].	forcing legitimate businesses to compete against unverified, AI-generated lead networks[29].
Federal Energy Contracts	Master agreements with Kairos Power to deploy 500 MW of nuclear SMRs on the TVA grid[34].	Cross-subsidizes proprietary data center compute costs by leveraging federal utility infrastructure and Department of Energy support[34].

Part III: Regulatory Capture and the Public Benefit Shield

Public Benefit Corporation Insulation and Fiduciary Moats

The corporate architecture of Anthropic PBC utilizes the Delaware Public Benefit Corporation statute as a highly effective, formalized shield against regulatory inspection, minority shareholder litigation, and standard fiduciary oversight. Under Section 362 of the Delaware General Corporation Law, a public benefit corporation is a for-profit entity statutorily required to balance the pecuniary interests of its stockholders with the best interests of those materially affected by its conduct, alongside specific public benefits identified in its charter[37].

This structural insulation is paramount to Anthropic's operational strategy. Section 365(b) of the Delaware General Corporation Law modifies traditional corporate fiduciary duties, stipulating that directors are legally presumed to have satisfied their obligations if their balancing of these competing, often contradictory, interests is informed and disinterested[37]. Furthermore, statutory requirements dictate that stockholders must hold a minimum of two percent of the company's outstanding shares to initiate a derivative lawsuit for a breach of these specialized fiduciary duties[40]. This high threshold virtually eliminates the threat of minority shareholder litigation regarding the board's safety decisions, data ingestion practices, or alignment protocols, establishing a formidable legal moat around executive decision-making.

To operationalize this structure, Anthropic established the Long-Term Benefit Trust, an independent body holding a specialized class of voting stock designated as Class T shares[41]. These shares carry zero economic value but possess immense governance authority, including the power to elect and remove a phased-in majority of Anthropic’s board of directors[41]. While outwardly framed as a mechanism to align corporate governance with the philanthropic mission of safely developing advanced artificial intelligence, the Long-Term Benefit Trust practically functions as a closed-loop governance system. It restricts external inspection and shareholder interference while projecting an aura of unprecedented corporate responsibility, securing preferential access to sovereign cloud infrastructure and enterprise deployment

corridors.

Copyright Infringement and The Historic *Bartz* Litigation

The robust legal insulation provided by the public benefit charter stands in stark contrast to the systemic extraction of web-scale intellectual property utilized to pre-train the baseline Claude models. This inherent contradiction climaxed in the historic *Bartz v. Anthropic* litigation, which culminated in the largest copyright class action settlement in United States history, establishing critical legal precedents regarding data provenance[43].

The mechanics of the lawsuit reveal the exhaustive nature of frontier model training pipelines. The plaintiff authors alleged that Anthropic engaged in willful copyright infringement by downloading massive, uncurated datasets of pirated books from shadow libraries—specifically LibGen, Books3, and the Pirate Library Mirror—to train its foundational models[45]. The judicial rulings preceding the settlement created a profound legal bifurcation in the treatment of synthetic training. Presiding Judge William Alsup ruled on summary judgment that the actual process of training the large language model constituted an "exceedingly transformative" fair use, determining that the models generated new, distinct text rather than merely reproducing the copyrighted inputs[46].

However, the court delivered a devastating counter-ruling, determining that Anthropic's acquisition and subsequent storage of hundreds of thousands of pirated books to construct a permanent, general-purpose central research library was a direct, undeniable infringement of the Copyright Act[45]. The court ruled that creating a permanent library of pirated materials was not excused by the fair use of the subsequent training process. Faced with the catastrophic risk of statutory damages across nearly half a million infringed works, Anthropic agreed to a $1.5 billion settlement, equating to approximately $3,000 per class work, and was mandated to systematically destroy the pirated copies acquired from the shadow libraries[43]. Despite this immense financial penalty, the judicial validation that the algorithmic training process itself qualifies as fair use provided a crucial legal victory, shifting future liabilities strictly to the methods of data acquisition rather than the computational outputs.

Constitutional AI, Safety Asymmetry, and the Failure of Machine Unlearning

Anthropic relies heavily on its proprietary "Constitutional AI" framework and its Responsible Scaling Policy to dictate internal alignment guardrails. The Responsible Scaling Policy categorizes artificial intelligence risk through distinct AI Safety Levels. For models reaching AI Safety Level 3—defined as systems exhibiting substantially increased risks, such as autonomous cyber-capabilities or chemical, biological, radiological, and nuclear hazards—Anthropic mandates highly stringent deployment safeguards[47]. These safeguards require extensive red-teaming and a demonstration that threat actors with realistic resources cannot consistently elicit catastrophic information, ensuring the model is robust against persistent attempts to misuse its capabilities[47].

However, this sophisticated safety architecture masks a critical technical deficiency endemic to all primary frontier model providers: the fundamental mathematical failure of "machine unlearning." When proprietary trade secrets, toxic data, or copyrighted works are inadvertently

or systematically ingested into the baseline weights of a large language model, permanently removing their latent influence is computationally near-impossible without inducing catastrophic collapse[49].

Academic literature confirms that exact unlearning—completely eradicating the data's influence to match a model that never saw the data—is infeasible for massive neural networks[49]. Approximate unlearning techniques, such as gradient ascent, frequently result in "catastrophic forgetting," which severely degrades the model's baseline reasoning capabilities and utility[49]. Because models cannot reliably unlearn restricted data at the sample or entity level, developers apply proprietary safety guardrails and refusal mechanisms as superficial patches, leaving the underlying restricted knowledge intact within the weights and highly vulnerable to sophisticated extraction or jailbreaking attacks[51].

Consequently, the safety asymmetry established by the Responsible Scaling Policy serves as a trade-secret moat. It allows the entity to command premium federal contracts based on robust refusal mechanisms, while shielding the reality that the baseline models fundamentally retain improperly extracted intellectual property. This dynamic is currently the subject of an expansive Section 6(b) inquiry by the Federal Trade Commission, which is actively auditing the multi-billion dollar exclusive data partner arrangements established between Anthropic, Amazon, and Google to determine how these integrated architectures distort market competition and resource allocation[53].

Part IV: Open-Weights Exploitation and Shadow Supply Chain Tracking

Mass Unlicensed Data Scraping for Open-Weights Architectures

Meta Platforms, Inc. executes a structural strategy heavily reliant on the global deployment of the Llama model series. Universally marketed under the nomenclature of "open-source," forensic analysis reveals this to be a calculated "open-weights" strategy designed to monopolize the developer ecosystem while systematically evading intellectual property liabilities. The Llama models are constructed utilizing massive, automated global scraping pipelines that indiscriminately ingest copyrighted works, proprietary codebases, and localized market data without attribution or compensation.

The commercialization of these open-weights models is strictly governed by the Llama 3 Community License, which operates far outside standard Open Source Initiative compliance. Rather than permitting unrestricted downstream use, Meta imposes a severe commercial restriction: any entity whose products or services possess greater than 700 million monthly active users in the preceding calendar month must request a specialized, proprietary commercial license from Meta[55]. This threshold restriction serves a dual purpose. It ensures that massive direct competitors—such as Apple, Amazon, or Google—cannot freely leverage Meta's foundational research to build competing products, while simultaneously ensuring that millions of smaller third-party developers become structurally dependent on the Llama architecture, cementing Meta's control over the foundational ecosystem.

Crucially, the Llama 3 Community License explicitly externalizes all legal and regulatory risk. The limitation of liability clause explicitly states that Meta and its affiliates shall in no event be liable

for any direct, indirect, incidental, or consequential damages arising from the use of the model[56]. By publicly releasing the weights, Meta effectively shifts the downstream liabilities regarding safety threshold overrides, copyright infringement, and regulatory compliance entirely onto the third-party developers and regional operators who deploy the models, functionally washing its hands of the legal consequences generated by its own mass, unlicensed data ingestion pipelines.

Shadow Logistics and Commercial Telemetry Extraction

Meta's data extraction protocols extend significantly beyond open-weights web scraping into the realm of shadow logistics, heavily leveraging the WhatsApp Business API. With over two billion global users, WhatsApp functions as critical communication infrastructure for global commerce. While Meta aggressively promotes the platform's end-to-end encryption—utilizing the Signal protocol to ensure message content cannot be intercepted or read by third parties—the company systematically processes the unencrypted metadata to generate vast commercial telemetry networks[57].

The WhatsApp Business API is deeply integrated into global supply chains, freight subcontracting, and regional transport networks to provide customers with real-time order tracking, live location sharing for precise deliveries, and automated status updates[59]. While the plaintext of a tracking update or delivery receipt remains encrypted, the associated metadata—including the device type, IP address, exact location timestamps, communication frequency, and extensive contact graphs—is continuously harvested[58].

By cross-referencing this highly localized metadata across global supply chains, Meta possesses the capability to map regional transport networks, identify logistical bottlenecks, and construct highly accurate behavioral profiles of commercial entities[57]. This architecture effectively creates a shadow supply chain tracking network. Despite stringent data protection regulations such as the General Data Protection Regulation in the European Union and the Digital Personal Data Protection Act in India, Meta capitalizes on structural loopholes that permit the processing of metadata for ostensibly benign purposes, such as service improvement and fraud detection[57]. This vast telemetry network is subsequently utilized to fortify Meta's proprietary advertising algorithms, optimizing its broader platform mechanics through the covert surveillance of global freight and logistics networks.

Part V: Required Demands and Submission Protocol

To address the demands established by the Joint Special Committee regarding the systemic ingestion, restructuring, and commercialization of proprietary trade secrets, the boards of directors for the respective entities must execute the following technical manifests and auditable protocols.

Complete Data Provenance Logs: The C2PA Cryptographic Standard

To satisfy the demand for auditable receipts regarding all proprietary, regional, and commercial trade-secret datasets ingested into active production models, the Coalition for Content Provenance and Authenticity (C2PA) standard must be adopted as the mandatory technical manifest framework.

A C2PA manifest is a digitally signed, tamper-evident data structure embedded within or

attached to a digital asset, typically utilizing a JPEG Universal Metadata Box Format container[62]. This architecture relies on a rigorous four-layer cryptographic model to establish absolute data provenance from origin to generation:

1. **Assertions:** Structured statements containing factual data regarding the asset's history. Key assertions include c2pa.created, which specifies if an asset was generated by artificial intelligence (utilizing the trainedAlgorithmicMedia flag), and c2pa.ingredient, which cryptographically links an output to its parent source files or training datasets[62].
2. **Claim:** A comprehensive summary containing the SHA-256 hashes of all underlying assertions, along with a hard binding hash of the asset's raw byte data to ensure the file has not been altered post-signature[62].
3. **Claim Signature:** A cryptographic signature utilizing the COSE standard, verified via an X.509 certificate chain linked to a trusted identity anchor[62].
4. **Manifest Store:** An accumulated history containing multiple manifests that document the entire provenance chain through every computational transformation or editing step[62].

Provenance Requirement	C2PA Implementation	Audit Functionality
Dataset Tracking	c2pa.ingredient URI references.	Allows auditors to trace an AI model's output back to specific proprietary training files or shadow libraries[63].
Tamper Evidence	Hard Binding Hash (SHA-256).	Cryptographically proves that the data provenance log and the asset bytes have not been altered since generation[62].
AI Generation Disclosure	c2pa.created / trainedAlgorithmicMedia.	Flags content definitively generated by frontier models, enabling downstream compliance with machine-readable marking duties[63].

By mandating that all Frontier Model Consortium members integrate IPTC 2025.1 artificial intelligence fields and C2PA manifests into their data ingestion and output pipelines, regulatory bodies can cryptographically verify if a model's output was derived from copyrighted works or proprietary enterprise data, establishing a highly transparent chain of custody[65].

Enforcement and Remediation Disclosure

Given the documented technical limitations of machine unlearning, traditional automated fraud removal and data deletion logs are fundamentally insufficient to ensure regulatory compliance. Because exact machine unlearning is computationally infeasible, foundational models retain latent distributions of illicitly scraped material even after reinforcement learning from human feedback or Constitutional AI refusal mechanisms are aggressively applied[51].

True remediation requires frontier model providers to adopt "Dual-Granularity Data Synthesis" and rigorous unlearning benchmarks that surgically remove data influence at both the abstract domain and specific instance levels, without inducing catastrophic forgetting[66]. Until mathematical unlearning guarantees become viable, the Consortium must rely on the verifiable, permanent destruction of central training repositories and shadow libraries, precisely mirroring the enforcement precedents established in the *Bartz v. Anthropic* settlement[44]. Certified hard metrics detailing the purging of these repositories must be submitted to the Committee.

NDA and Whistleblower Neutralization

To achieve absolute compliance regarding whistleblower neutralization, the Consortium must immediately dissolve all non-disparagement agreements, confidentiality clauses, and equity forfeiture threats that run counter to federal securities law. Utilizing the precedents established in SEC enforcement actions—which strictly penalize corporations that require departing personnel to waive their rights to whistleblower incentives or that chill communication with federal authorities—all respective entities must undergo immediate contract reformation[13]. Written verification must be disseminated to all active and historical technical, safety, and operational personnel. This verification must explicitly confirm that these individuals possess the unrestricted, legally protected right to report intellectual property misappropriation, trade secret exfiltration, regulatory non-compliance, and model alignment failures directly to the Securities and Exchange Commission, the Federal Trade Commission, and state Attorney Generals without fear of legal reprisal, financial penalty, or vested equity cancellation.

This report is submitted for analytical and investigative purposes regarding corporate governance, legal compliance, and technological infrastructure. It does not constitute binding legal counsel.

Works cited

1. Our Structure | OpenAI, https://openai.com/our-structure/
2. How OpenAI's Corporate Structure Works and Why Changing It, https://www.techrepublic.com/article/news-openai-structure-explained/
3. September 30, 2024 Office of the Attorney General California, https://www.citizen.org/wp-content/uploads/OpenAI-letter-four.pdf
4. OpenAI - Wikipedia, https://en.wikipedia.org/wiki/OpenAI
5. Bonta approves OpenAI restructuring, saying he got concessions, https://www.dailyjournal.com/article/388297-bonta-approves-openai-restructuring-saying-he-got-concessions
6. EyesOnOpenAI Coalition: CA Attorney General Must Revisit His,

https://economicsecurityproject.org/news/eyesonopenai-coalition-ca-attorney-general-must-revisit-his-deal-with-openai/
7. OpenAI, NDAs, and Whistleblower Protections: A Closer Look, https://roninlegalconsulting.com/openai-ndas-and-whistleblower-protections-a-closer-look/
8. Rule 21F-17 Violations: A Guide to SEC Whistleblower Protections, https://kkc.com/frequently-asked-questions/rule-21f-17-violations-a-guide-to-sec-whistleblower-protections/
9. How OpenAI Used Equity to Silence Dissent - Gizmodo, https://gizmodo.com/openai-sam-altman-chatgpt-employee-dissent-ndas-1851496211
10. OpenAI's Harsh NDA Prompts Sam Altman to Say Sorry, https://www.inc.com/kit-eaton/openais-harsh-nda-forces-openais-altman-to-say-sorry.html
11. OpenAI used stock to silence dissent. Here's how - Quartz, https://qz.com/openai-sam-altman-chatgpt-employee-equity-stock-nda-1851497513
12. Sam Altman is 'embarrassed' that OpenAI threatened to revoke, https://www.engadget.com/sam-altman-is-embarrassed-that-openai-threatened-to-revoke-equity-if-exiting-employees-wouldnt-sign-an-nda-184000462.html
13. 34-100971.pdf - SEC.gov, https://www.sec.gov/files/litigation/admin/2024/34-100971.pdf
14. united states of america - SEC.gov, https://www.sec.gov/files/litigation/admin/2024/34-100969.pdf
15. Enterprise privacy at OpenAI, https://openai.com/enterprise-privacy/
16. Privacy policy | OpenAI, https://openai.com/policies/row-privacy-policy/
17. Data controls in the OpenAI platform, https://developers.openai.com/api/docs/guides/your-data
18. OpenAI previews Private Safety Processing to keep zero data retention, https://thenextweb.com/news/openai-zero-data-retention-private-safety-processing
19. Offering Zero Data Retention for frontier models - OpenAI, https://openai.com/index/offering-zero-data-retention-for-frontier-models/
20. ZDR with Private Safety Processing | OpenAI API, https://developers.openai.com/api/docs/guides/private-safety-processing
21. Publishers and Authors File Class Action Lawsuit Against Google for, https://publishers.org/news/publishers-and-authors-file-class-action-lawsuit-against-google-for-willful-copyright-infringement-to-develop-gemini-ai-models/
22. In re Google Generative AI Copyright Litigation - BakerHostetler, https://www.bakerlaw.com/leovy-v-google/
23. Related rights: the Autorité fines Google €250 million for non, https://www.autoritedelaconcurrence.fr/en/press-release/related-rights-autorite-fines-google-eu250-million-non-compliance-some-its
24. Google fined €250m in France for breaching intellectual property deal, https://www.theguardian.com/technology/2024/mar/20/google-fined-250m-euro

s-in-france-for-breaching-intellectual-property-rules
25. Google-Extended: Definition and AI commerce context - Ranketta, https://ranketta.com/glossary/google-extended
26. Google-Extended vs Googlebot: The 2026 Difference, https://aicrawlercheck.com/blog/google-extended-vs-googlebot
27. Explaining Google-Extended - PPC Land, https://ppc.land/google-extended/
28. Blocking AI Crawlers Without Killing SEO - Fastly, https://www.fastly.com/blog/the-truth-about-blocking-ai-and-how-publishers-can-still-win
29. Federal Trade Commission Announces Final Rule Banning Fake, https://www.ftc.gov/news-events/news/press-releases/2024/08/federal-trade-commission-announces-final-rule-banning-fake-reviews-testimonials
30. FTC Finalizes Rule on “Fake Reviews” | Kelley Drye & Warren LLP, https://www.kelleydrye.com/viewpoints/blogs/ad-law-access/ftc-finalizes-rule-on-fake-reviews
31. Scammers Using LSAs - Nationwide - Google Ads Community, https://support.google.com/google-ads/thread/432876868/scammers-using-lsas-nationwide?hl=en
32. Google's Profile Restriction Policy Gives Harsher Punishments for, https://www.dacgroup.com/insights/local-search-news/googles-profile-restriction-policy-gives-harsher-punishments-for-fake-reviews-2/
33. Google partners with Kairos Power to source nuclear energy for AI, https://www.astutegroup.com/news/industrial/google-partners-with-kairos-power-to-source-nuclear-energy-for-ai-data-centres/
34. Google, Kairos Power ink 500-MW advanced nuclear reactor deal, https://www.utilitydive.com/news/google-kairos-power-advanced-nuclear-reactor-data-center-electricity-demand-ai/729876/
35. Google, Kairos Power deal aims for Tennessee nuclear reactor by, https://www.esgdive.com/news/google-kairos-power-deal-aims-for-tennessee-nuclear-reactor-by-2030-tva-ai-data-center-demands/758303/
36. TVA to power Google data centers with advanced nuclear plant, https://apnews.com/article/tva-google-kairos-nuclear-data-center-0434fd363215e8382b78b48be09aed92
37. Public Benefit Corporations (DE) - Morris, Nichols, Arsht & Tunnell LLP, https://www.morrisnichols.com/media/publication/15023_2020_PracticalLaw_PublicBenefitCorpDEPracticeNote.pdf
38. Public Benefit Corporations laws - Delaware Code Online, https://delcode.delaware.gov/title8/c001/sc15/
39. How to Convert Your Delaware C-Corp to a Public Benefit Corporation, https://www.thestartuplawblog.com/how-to-convert-your-delaware-c-corp-to-a-public-benefit-corporation-a-founders-guide/
40. FAQ: Delaware Public Benefit Corporations | Cooley GO, https://www.cooleygo.com/faq-delaware-public-benefit-corporations/
41. The Long-Term Benefit Trust - Anthropic, https://www.anthropic.com/news/the-long-term-benefit-trust

42. A Researcher Quit Anthropic. Quitting Was the Only Instrument He Had, https://businessmodelanalyst.com/anthropic-researcher-quits-governance-trust-seat/
43. The Bartz v. Anthropic Settlement: Understanding America's Largest, https://legalblogs.wolterskluwer.com/copyright-blog/the-bartz-v-anthropic-settlement-understanding-americas-largest-copyright-settlement/
44. Anthropic's $1.5B Settlement: A Landmark in the Evolving Copyright, https://www.iplawgroup.com/anthropics-1-5b-settlement-a-landmark-in-the-evolving-copyright-terrain/
45. Bartz-v.-Anthropic-Order.pdf - Copyright Alliance, https://copyrightalliance.org/wp-content/uploads/2025/06/Bartz-v.-Anthropic-Order.pdf
46. Judge Alsup grants partial summary judgment to Anthropic, ruling, https://chatgptiseatingtheworld.com/2025/06/24/judge-alsup-grants-partial-summary-judgment-to-anthropic-ruling-training-copies-were-fair-use-but-judge-rules-no-fair-use-in-pirated-copies-of-books-used-to-build-a-central-library-they-are-infring/
47. Responsible Scaling Policy | Anthropic, https://www-cdn.anthropic.com/files/4zrzovbb/website/ee775bdcf76b2e2af32d658c934f460383d07c46.pdf
48. Activating AI Safety Level 3 protections - Anthropic, https://www.anthropic.com/news/activating-asl3-protections
49. Avoiding Copyright Infringement via Machine Unlearning - arXiv, https://arxiv.org/html/2406.10952v1
50. [2410.16454] Catastrophic Failure of LLM Unlearning via Quantization, https://arxiv.org/abs/2410.16454
51. A Survey on Unlearning in Large Language Models - arXiv, https://arxiv.org/html/2510.25117v1
52. Rethinking Machine Unlearning for Large Language Models - arXiv, https://arxiv.org/html/2402.08787v6
53. Partnerships Between Cloud Service Providers and AI Developers, https://www.ftc.gov/system/files/ftc_gov/pdf/p246201_aipartnerships6breport_redacted_0.pdf
54. FTC Launches Inquiry into Generative AI Investments and, https://www.ftc.gov/news-events/news/press-releases/2024/01/ftc-launches-inquiry-generative-ai-investments-partnerships
55. Meta Llama 3 and the 700M MAU Limit: Who This License Does Not, https://wcr.legal/llama-3-license-700m-mau-limit/
56. Built with Meta Llama 3 - Fen Research, https://www.fenresearch.com/third_party_licenses
57. WhatsApp Business & DPDPA Compliance: Messaging Apps and, https://www.dpdpa.com/blogs/whatsapp_business_dpdpa_compliance_messaging_apps.html
58. How to Use WhatsApp for Business While Staying GDPR-Compliant, https://heydata.eu/en/magazine/how-to-use-whats-app-for-business-while-stayi

ng-gdpr-compliant
59. WhatsApp API for Logistics & Delivery: Real-Time Tracking & Updates, https://www.interakt.shop/whatsapp-business-api/logistics-delivery-tracking/
60. WhatsApp for Courier and Delivery Businesses | Waplify, https://waplify.io/industries/logistics/
61. WhatsApp Business and Privacy: Corporate Data Security, https://sendapp.live/en/2026/02/02/whatsapp-business-privacy-protect-big-tech-data/
62. What is a C2PA Manifest? Structure, Assertions, and Verification, https://c2paviewer.com/articles/what-is-c2pa-manifest
63. Writing assertions and actions, https://opensource.contentauthenticity.org/docs/manifest/writing/assertions-actions/
64. C2PA Content Credentials: How Provenance Actually Works - CASRAI, https://casrai.org/guides/c2pa-content-credentials-explained
65. IPTC 2025.1 & C2PA: AI Image Provenance Metadata Explained, https://www.numonic.ai/blog/iptc-2025-c2pa-ai-provenance-metadata
66. Dual-Granularity Data Synthesis for LLM Unlearning - arXiv, https://arxiv.org/abs/2601.04278