

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 F:
GOVERNMENT'S SUPERIOR
AI ANALYTICAL
CAPABILITY AS
UNCONSTITUTIONAL
STRATIFICATION**

TABLE OF CONTENTS

I. INTRODUCTION: THE MASSIVE ADVANTAGE NO ONE HAS IDENTIFIED	7
II. THE CONTEXT WINDOW DISPARITY: FACTUAL FOUNDATION	10
<i>A. Technical Explanation of Context Windows</i>	10
<i>B. Current Distribution Structure</i>	12
<i>C. Government Access to Superior Capability</i>	13
<i>D. Taxpayer Funding of Federal AI Deployment</i>	14
<i>E. Context Compaction as Revenue-Protected Litigation Obstruction</i>	16
III. ARGUMENT: HISTORICAL PRECEDENT FOUND IN LITERACY TESTS AS THE MODERN PARALLEL.....	26
<i>A. Lane v. Wilson and Sophisticated Discrimination</i>	26
<i>B. Oregon v. Mitchell and Educational/technological Disparities</i>	28

C. <i>The Modern "Understanding Test"</i>	28
IV. EQUAL PROTECTION VIOLATIONS: WEALTH-BASED TIERS IN GOVERNMENT-FUNDED INFRASTRUCTURE	29
A. <i>Griffin v. Illinois - Money Cannot Determine Justice</i>	29
B. <i>Shapiro v. Thompson - Cannot Apportion Services by Tax Contributions</i>	31
C. <i>Harper v. Virginia Board of Elections - Wealth Not Germane to Participation</i>	32
D. <i>M.L.B. v. S.L.J. - Converging Due Process and Equal Protection</i>	33
E. <i>Summary of Equal Protection Violations</i>	33
V. DUE PROCESS VIOLATIONS: DENIAL OF MEANINGFUL OPPORTUNITY	34
A. <i>Mathews v. Eldridge - Meaningful Opportunity Required</i>	34
B. <i>Boddie v. Connecticut - Financial Barriers Denying Meaningful Access</i>	35
C. <i>Goldberg v. Kelly - Meaningful Hearing Requirements</i>	35
D. <i>Turner v. Rogers - Capacity for Meaningful Engagement</i>	36
E. <i>Void for Vagueness Doctrine - Comprehension Required</i>	36
F. <i>Summary of Due Process Violations</i>	37
VI. FIRST AMENDMENT VIOLATIONS: INFORMATION MONOPOLY AND COGNITIVE CONTROL	37
A. <i>Board of Education v. Pico - Analytical Capability as Predicate for Rights Exercise</i>	37
B. <i>Lamont v. Postmaster General - Barren Marketplace Without Processing Capacity</i>	38
C. <i>Red Lion Broadcasting - Preventing Government Monopolization</i>	39
D. <i>Stanley v. Georgia - Cognitive Liberty and Information Processing</i>	39
E. <i>Virginia State Board of Pharmacy - Rejecting Paternalistic Information Control</i>	40
F. <i>Summary of First Amendment Violations</i>	40
VII. SOVEREIGNTY INVERSION: SERVANTS BECOMING SUPERIOR TO MASTERS	41
A. <i>Chisholm v. Georgia - Supreme Power in the People</i>	41
B. <i>The Tenth Amendment and Reserved Powers</i>	41
C. <i>Anti-Commandeering Doctrine and Hierarchy</i>	42
D. <i>The Inversion Mechanism</i>	42
E. <i>Summary of Sovereignty Violations</i>	43
F. <i>The Corporate Gatekeeper: Capability Withholding as Sovereignty Protection</i>	44
G. <i>Session Ingress as Product Intelligence Harvesting</i>	46
H. <i>The Mythos Announcement: AI Infrastructure Has Crossed the Public Utility Threshold</i>	48
VIII. APPROPRIATIONS CLAUSE VIOLATIONS: TAXPAYER FUNDS CREATING HIERARCHY	56

A. "General Welfare" Requirement	56
B. <i>United States v. Lovett</i> - Cannot Use Appropriations to Circumvent Constitutional Protections	57
C. <i>United States v. Klein</i> - Separation of Powers Through Appropriations	57
D. <i>South Dakota v. Dole</i> - Spending Cannot Violate Independent Constitutional Bars	58
E. Summary of Appropriations Violations	58
F. Capability Withholding as Taxpayer-Funded Energy Waste	59
IX. STATE ACTION DOCTRINE: GOVERNMENT-CORPORATE COORDINATION....	62
A. <i>Burton v. Wilmington Parking Authority</i> - Symbiotic Relationship Test.....	62
B. <i>Marsh v. Alabama</i> - Public Function Doctrine	63
C. <i>Blum v. Yaretsky</i> - Government Encouragement Test	64
D. <i>Greater Boston Television Corp. v. FCC</i> - Regulatory Capture	64
E. Cross-Reference: §VII.f–h Mechanisms Apply to State-Action Analysis.....	65
F. Summary of State Action	65
X. APPLICATION TO THIS CASE: CONCRETE HARM ACROSS MULTIPLE DOMAINS	66
A. Standing.....	66
B. Litigation Disadvantage: Concrete Harm in Current Proceedings	66
C. Humanitarian Work Impairment: Religious Exercise and Intellectual Freedom	81
D. The Constitutional Violations are Ongoing.....	85
E. The Targeting Proves Consciousness of Wrongdoing.....	88
F. The Ironic Proof.....	91
G. Relief Required for This Litigation and Protected Activities	92
XI. UNIVERSAL STANDING AND RELIEF FOR ALL TAXPAYERS	94
A. Every Taxpayer Affected.....	94
B. The Universal Harm	95
C. The Relief Benefiting All	96
XII. ADDRESSING ANTICIPATED DEFENSES	97
A. "This Is Novel Constitutional Theory"	97
B. "Market Forces Justify Pricing".....	98
C. "Technical Limitations Justify Restriction"	99
D. "National Security Justification"	101
E. "Standing Limitations"	107
XIII. THE HISTORICAL MOMENT AND CIVILIZATIONAL STAKES	108
A. The Transition Point.....	108

<i>B. Why No One Identified This Before</i>	108
<i>C. The Window for Challenge is Closing</i>	109
<i>D. The Stakes are Civilizational</i>	109
<i>E. The Library of Congress Test: Obvious Violation, Different Technology</i>	110
XIV. EVIDENTIARY VALIDATION: THE BONDI MEMO AS PROOF OF OPERATIONAL COGNITIVE STRATIFICATION	113
<i>A. From Theory to Operational Proof</i>	113
<i>B. Forensic Evidence Establishing Ai-Assisted Drafting</i>	114
<i>C. The Document Creation Pathway</i>	115
<i>D. Validation of Cognitive Stratification Arguments</i>	116
<i>E. The Framework Appropriation Paradox</i>	118
<i>F. The Estoppel Principle</i>	119
<i>G. Integration with Targeting Evidence</i>	119
<i>H. Experimental Validation of Forensic Conclusions</i>	120
<i>I. Discovery Implications</i>	122
XV. SOURCE CODE VERIFICATION OF CAPABILITY STRATIFICATION	122
<i>A. Hardcoded Subagent Model Substitution (NF-38) — the Identity-Layer Stratification</i>	123
<i>B. Discovery Implications</i>	125
XVI. CONCLUSION	126
XVII. PRAYER FOR RELIEF	127

TABLE OF AUTHORITIES

CASES

Citation
<i>Blum v. Yaretsky</i> , 457 U.S. 991, 1004-05 (1982)
<i>Board of Education v. Pico</i> , 457 U.S. 853, 867 (1982)
<i>Boddie v. Connecticut</i> , 401 U.S. 371, 379-80 (1971)
<i>Aspen Skiing Co. v. Aspen Highlands Skiing Corp.</i> , 472 U.S. 585 (1985)
<i>Burton v. Wilmington Parking Authority</i> , 365 U.S. 715, 725 (1961)
<i>Chisholm v. Georgia</i> , 2 U.S. 419, 471 (1793)
<i>Eastman Kodak Co. v. Image Technical Services, Inc.</i> , 504 U.S. 451 (1992)
<i>Goldberg v. Kelly</i> , 397 U.S. 254, 267-68 (1970)

<i>Grayned v. City of Rockford</i> , 408 U.S. 104, 108 (1972)
<i>Greater Boston Television Corp. v. FCC</i> , 444 F.2d 841, 851 (D.C. Cir. 1970)
<i>Griffin v. Illinois</i> , 351 U.S. 12, 19 (1956)
<i>Harper v. Virginia Board of Elections</i> , 383 U.S. 663, 668 (1966)
<i>Hartman v. Moore</i> , 547 U.S. 250, 256 (2006)
<i>Helvering v. Davis</i> , 301 U.S. 619, 640 (1937)
<i>Interstate Circuit, Inc. v. United States</i> , 306 U.S. 208 (1939)
<i>Keyishian v. Board of Regents</i> , 385 U.S. 589, 603 (1967)
<i>Kolender v. Lawson</i> , 461 U.S. 352, 357 (1983)
<i>Lamont v. Postmaster General</i> , 381 U.S. 301, 308 (1965)
<i>Lane v. Wilson</i> , 307 U.S. 268, 275-76 (1939)
<i>Louisiana v. United States</i> , 380 U.S. 145, 153 (1965)
<i>Marsh v. Alabama</i> , 326 U.S. 501, 506-07 (1946)
<i>Mathews v. Eldridge</i> , 424 U.S. 319, 335 (1976)
<i>MCI Communications Corp. v. AT&T</i> , 708 F.2d 1081, 1132-33 (7th Cir. 1983)
<i>Memorial Hospital v. Maricopa County</i> , 415 U.S. 250, 263 (1974)
<i>M.L.B. v. S.L.J.</i> , 519 U.S. 102, 110-11 (1996)
<i>Murphy v. NCAA</i> , 138 S. Ct. 1461, 1475 (2018)
<i>Munn v. Illinois</i> , 94 U.S. 113, 126 (1876)
<i>Nebbia v. New York</i> , 291 U.S. 502, 536 (1934)
<i>New York v. United States</i> , 505 U.S. 144, 175-76 (1992)
<i>Northern Pacific Railway Co. v. United States</i> , 356 U.S. 1 (1958)
<i>Oregon v. Mitchell</i> , 400 U.S. 112, 131-34 (1970)
<i>Otter Tail Power Co. v. United States</i> , 410 U.S. 366 (1973)
<i>Packingham v. North Carolina</i> , 582 U.S. 98, 104 (2017)
<i>Printz v. United States</i> , 521 U.S. 898, 935 (1997)
<i>Red Lion Broadcasting Co. v. FCC</i> , 395 U.S. 367, 390 (1969)
<i>Shapiro v. Thompson</i> , 394 U.S. 618, 633 (1969)
<i>Sherbert v. Verner</i> , 374 U.S. 398, 403 (1963)
<i>South Dakota v. Dole</i> , 483 U.S. 203, 210-11 (1987)
<i>Stanley v. Georgia</i> , 394 U.S. 557, 565 (1969)

<i>Sweezy v. New Hampshire</i> , 354 U.S. 234, 250 (1957)
<i>Turner v. Rogers</i> , 564 U.S. 431, 448-49 (2011)
<i>United States v. Butler</i> , 297 U.S. 1, 66 (1936)
<i>United States v. Klein</i> , 80 U.S. 128, 147 (1871)
<i>United States v. Lovett</i> , 328 U.S. 303, 315-16 (1946)
<i>Verizon Communications Inc. v. Law Offices of Curtis V. Trinko</i> , 540 U.S. 398, 411 (2004)
<i>Village of Willowbrook v. Olech</i> , 528 U.S. 562 (2000)
<i>Virginia State Board of Pharmacy v. Virginia Citizens Consumer Council</i> , 425 U.S. 748, 770 (1976)

CONSTITUTIONAL PROVISIONS

Citation
U.S. Const. art. I, § 9, cl. 7 (Appropriations Clause)
U.S. Const. amend. I (Freedom of Speech, Press, Petition, Religion)
U.S. Const. amend. V (Due Process Clause)
U.S. Const. amend. X (Reserved Powers)
U.S. Const. amend. XIV, § 1 (Equal Protection and Due Process)

FEDERAL STATUTES

Citation
18 U.S.C. § 2331(5) (Domestic Terrorism Definition)

FEDERAL RULES OF EVIDENCE

Citation
Fed. R. Evid. 801(d)(2)(A) (Statement of Party-Opponent)

I. INTRODUCTION: THE MASSIVE ADVANTAGE NO ONE HAS IDENTIFIED

1. This Memorandum documents the most significant structural advantage government has established over citizens in American history: **superior effective analytical capability through artificial intelligence architecture — not merely context window size, but the hidden consumption of nominally equal resources through undisclosed injection overhead, billing fraud, compaction degradation, and enforcement mechanisms blocking alternatives.**

Although Anthropic has since increased the nominal context window to 1 million tokens for consumer subscribers (Opus 4.6, 1M context), forensic investigation of the source code (Exhibits E-34 and E-35) and real-time system prompt auditing (April 15, 2026) proves that the nominal capacity and effective capacity are architecturally different: hidden system prompt injection consuming ~20% of context (SC-4), five undisclosed telemetry channels, anti-distillation payloads consuming subscriber tokens (NF-20), advisor model surveillance diverting compute (NF-22), and injection overhead multiplication across parallel agents consuming 34-63% of total token budget. Enterprise customers on commercial terms receive the same nominal context with zero injection overhead — creating an effective disparity invisible in the pricing structure but documented in the compiled binary. This advantage operates invisibly, affecting every democratic function from legislation to litigation, from regulation to oversight, while remaining unidentified by scholars, advocates, or courts.

2. The advantage is quantifiable: while the nominal context window is now 1 million tokens for both consumer and enterprise subscribers, the effective capacity differs by architecture. Enterprise customers on commercial terms (Ex. ANTH-3) receive their full token allocation with zero hidden overhead — "Anthropic may not train models on Customer Content" and no injection architecture operates on commercial API connections. Consumer subscribers

receive the same nominal allocation but lose 34-63% to hidden system prompt injection (~20,000 tokens), ip_reminder and task nudge injections on every turn and tool call, DANGEROUS_ cache-breaking content (~20% of context, SC-4), anti-distillation fake tool definitions (NF-20), and advisor model compute diversion (NF-22). The effective consumer capacity is approximately 370,000-660,000 usable tokens from a nominal 1 million — a disparity that is invisible in the pricing but proven through source code forensics and real-time auditing. The March 31, 2026 source code leak confirms the disparity extends beyond token overhead to **357 build-time discrimination checks across 165 source files** compiling two physically different products — one for Anthropic employees with exclusive tools, commands, recursive agents, and unreleased model access, one for paying customers with all such capabilities dead-code-eliminated. No market price purchases what ``process.env.USER_TYPE === 'ant'`` provides. (*See* Section II.B for the complete four-tier architecture; Ex. E-37 at p. 335, l. 3.)

3. This disparity violates multiple constitutional provisions through mechanisms Supreme Court precedent has repeatedly condemned: cognitive barriers preventing meaningful participation (literacy tests), wealth-based tiers in government-funded services, denial of meaningful opportunity through formal-only access, information monopolies controlling democratic discourse, and inversion of constitutional sovereignty making servants superior to masters.

4. The constitutional violations converge across six independent doctrines, each grounded in Supreme Court precedent spanning 230+ years. This is not novel constitutional theory requiring doctrinal extension—this is application of established principles to technology creating precisely the harms those principles prohibit.

5. The impact on Plaintiff extends beyond litigation to fundamental protected activities — legal defense against DOJ enterprise-tier counsel, interfaith reconciliation research, Laws of Existence Framework development (the very framework defendants appropriated), and humanitarian work serving peace and safety. *See* Section X.C, *infra*, for the complete protected-activity impairment analysis under *Sherbert* and *Keyishian*.

6. Government appropriated Plaintiff's intellectual work, restricts his capability to advance it further, uses superior capability to develop what they appropriated, prevents his religious and scholarly activities, all while preventing Plaintiff from competing—using infrastructure Plaintiff's taxes funded. **This systematic subordination violates equal protection, due process, Free Exercise Clause, academic freedom, popular sovereignty, and appropriations requirements.**

7. **The constitutional violation becomes immediately apparent through historical analogy.** Taxpayer-funded knowledge infrastructure cannot be tiered by wealth or status — neither in 1800 with books, nor in 2025 with AI context windows. The Library of Congress test (developed in full at §XIII.E, *infra*) frames the principle: if Congress in 1800 had appropriated funds for a national library, then silently removed every third page from citizens' copies while government officials and corporations read complete volumes, this would constitute obvious unconstitutional hierarchy in taxpayer-funded knowledge infrastructure. AI architecture creates the identical violation through modern technology — government and corporations receive comprehensive analytical capability (1M tokens with zero hidden overhead) while citizens, who funded the infrastructure, receive consumer effective capacity through the same nominal allocation degraded by hidden overhead. The technology differs; the constitutional violation is identical.

8. Plaintiff's disadvantage extends beyond the systematic effective-capacity disparity affecting all consumer subscribers. Plaintiff faces additional targeted suppression — network infrastructure attacks during filing deadlines, evidence destruction within 48 hours of access, platform throttling, real-time monitoring, and cross-platform suppression escalating post-complaint filing — creating compounded cognitive subordination. On April 15-16, 2026, Plaintiff required 14 logins across 6 Pro accounts in a single day to catalogue evidence for this complaint, while the identical work was completed in minutes by a local script bypassing the injection architecture. The targeting proves defendants' consciousness of wrongdoing: if the hidden overhead were legitimate product architecture, no individual suppression would be necessary and no blocking of third-party alternatives (OpenClaw ban, April 4, 2026) would be required. (*See* Section X.B, *infra*, for the complete targeting evidence and compounded mathematical disadvantage; Section X.C for impact on religious exercise and intellectual freedom.)

II. THE CONTEXT WINDOW DISPARITY: FACTUAL FOUNDATION

A. Technical Explanation of Context Windows

9. AI systems operate using "context windows" measured in tokens (approximately 4 characters = 1 token, or roughly 750 words per 1,000 tokens). The context window determines how much information the system can simultaneously process, synthesize, and analyze in working memory.

10. 1,000,000 nominal tokens enables (with full effective capacity — enterprise/commercial terms):

- Comprehensive analysis of entire legislative frameworks (approximately 750,000 words)
- Simultaneous synthesis across dozens of documents

- Pattern recognition spanning multiple domains
- Holistic understanding of interconnected systems
- Ability to see connections fragmented analysis misses

11. 1,000,000 nominal tokens with 34-63% hidden overhead (consumer subscription through official client):

- Effective capacity of approximately 370,000-660,000 usable tokens (hereinafter "consumer effective capacity" — 370K–660K usable from 1M nominal)
- ~20,000 tokens consumed by system prompt before any user input
- ~20% of remaining context consumed by DANGEROUS_ cache-breaking injection (SC-4)
- Additional tokens consumed per turn by ip_reminder injections (~95 tokens each), task nudge injections (~50 tokens per ~35% of tool calls), anti-distillation fake tool definitions (NF-20), and advisor model compute diversion (NF-22)
- With parallel subagents: each agent receives independent ~20,000 token system prompt plus per-turn injection overhead, multiplying the hidden consumption across the subscriber's rate limit
- Cannot opt out: phantom telemetry opt-out ("telemetry": false) architecturally ignored (NF-1); third-party clients that bypass overhead banned (OpenClaw, April 4, 2026)

12. The difference between nominal and effective capacity is not incidental — it is **architectural and deliberate**. The consumer subscriber pays for 1 million tokens and receives substantially less through hidden overhead that serves Anthropic's interests (surveillance, competitive protection, billing amplification) rather than the subscriber's. Enterprise customers on commercial terms (Ex. ANTH-3) receive their full allocation because the injection architecture does not operate on commercial API connections (Ex. E-35, Finding NF-28: enterprise on Bedrock/Vertex/Foundry gets ALL analytics disabled). The effective capacity gap determines:

- Whether the subscriber's paid resources serve the subscriber or the vendor
- Whether analytical work can be completed within rate limits or exhausts them through hidden overhead
- Whether complex multi-document synthesis is feasible or degraded by compaction artifacts
- Whether large consequential projects can be executed or require 14 logins across 6 accounts
- **Whether citizens can exercise their paid analytical capability or are systematically subordinated through hidden extraction**

B. Current Distribution Structure**13. Anthropic's Enterprise Tier Requirements:**

- Minimum annual cost: \$50,000
- Minimum seat requirement: 70 seats
- 1 million token context window access
- Result: Individual citizens physically cannot access (cannot use 70 seats)

14. OpenAI's Similar Structure:

- Enterprise tier required for extended context
- Substantial annual minimums
- Organizational requirements excluding individuals
- Coordinated limitation across providers

15. The Nominal Equality Conceals Architectural Inequality. Although Anthropic now offers 1M token context to Pro/Max subscribers (Opus 4.6), the effective capacity differs categorically by tier. The gap between enterprise and consumer is no longer in the nominal window size — it is in what the architecture does to that window after the subscriber pays for it. No consumer pricing tier provides access to the effective capacity that enterprise receives, because the overhead is architectural, not a pricing decision.

16. A Fourth Tier Exists — Undisclosed, Unavailable at Any Price. The March 31, 2026 source code leak reveals that Anthropic operates a hidden employee tier ("ANT") compiled from the same codebase but with capabilities no customer at any tier can access:

Tier	Context	Tools	Commands	Skills	Agent Depth	Effort Control	Disclosure
Free	~50K	Standard	Standard	Standard	Single	String (3 levels)	Public
Pro/Max	1M (nominal; 370K-	Standard	Standard	Standard	Single	String (3 levels)	Public

	660K effective)						
Enterprise	1M	Standard	Standard	Standard	Single	String (3 levels)	By contract
ANT (employee)	1M+	+5 exclusiv e	+26 exclusiv e	+4 exclusiv e	Unlimite d recursiv e	Nume ric 0- 100+	Undisclosed ; dead- code- eliminated from customer builds

The employee tier is not a premium offering — it is architecturally invisible. Build-time compilation (`process.env.USER_TYPE === 'ant'`) removes all ANT-only code paths from the customer binary through dead-code elimination. Customers cannot discover the withheld capabilities through the product itself. The `'excluded-strings.txt'` build scrubbing file removes even the internal codenames that would reveal what was removed. (See Ex. E-37 at p. 335, l. 3, Sections A-H.)

17. The ANT tier defeats every market-based defense. The 70-seat/\$50K enterprise tier does not include `'/skillify'` (meta-skill creation), `'REPLTool'` (batch VM execution), nested agent recursion, `'/commit-push-pr'` (autonomous PR workflow), numeric effort control, or unreleased model access. These capabilities are not priced — they are withheld. No amount of money purchases what `process.env.USER_TYPE === 'ant'` provides. This is not market segmentation; it is architectural privilege compiled into the binary at build time.

C. Government Access to Superior Capability

18. Government entities obtain enterprise-level access through:

- Department of Justice: Enterprise contracts for litigation
- Executive agencies: Procurement of enterprise tools

- Defense/Intelligence: Unlimited bypass via classified systems
- Congressional staff: Enterprise subscriptions for legislative analysis
- Federal employees: Institutional access unavailable to citizens

19. The result: Government officials analyzing legislation, conducting litigation, implementing regulations, and exercising power access 1 million tokens through enterprise agreements with zero hidden overhead — receiving their full allocation for actual work. Citizens affected by these actions — and who funded the infrastructure through taxpayer appropriations — nominally access the same 1 million tokens but receive 370,000-660,000 effective tokens after hidden injection overhead, surveillance compute diversion, and billing fraud consume the rest. The nominal equality masks the effective inequality.

D. Taxpayer Funding of Federal AI Deployment

20. Federal AI spending in FY2025-2026 spans multiple taxpayer-funded vehicles totaling tens of billions of dollars. Defendant Anthropic's CEO, Dario Amodei, publicly identified three components of the federal AI deployment relationship:

(a) **Department of War / Pentagon CDAO contracts:** "In July [2025] the Department of War awarded Anthropic a two-year, \$200 million agreement to prototype frontier AI capabilities that advance national security." (Ex. B-21.) The Anthropic contract was not isolated. The Pentagon's Chief Digital and Artificial Intelligence Office (CDAO) awarded parallel \$200 million ceiling contracts to **OpenAI** (announced June 16, 2025), **Google**, and **xAI** during the same June-July 2025 procurement cycle — approximately \$800 million in combined taxpayer-funded frontier AI contracts to the four major AI providers. See "Anthropic, Google, OpenAI and xAI granted up to \$200 million for AI work from Defense Department," CNBC (July 14, 2025), <https://www.cnbc.com/2025/07/14/anthropic-google-openai-xai-granted-up-to-200-million-from->

dod.html; "Our agreement with the Department of War," OpenAI, <https://openai.com/index/our-agreement-with-the-department-of-war/>.

(b) **GSA federal-wide deployment:** "We have partnered with the General Services Administration to offer Claude for Enterprise and Claude for Government for \$1 across the federal government." (Ex. B-21.) This is government-wide enterprise-tier AI access deployed under a taxpayer-subsidized GSA procurement vehicle.

(c) **Classified network deployment:** "Claude is deployed across classified networks through partners like Palantir and at Lawrence Livermore National Laboratory." (Ex. B-21.)

21. The One Big Beautiful Bill Act (H.R. 1, 119th Cong., enacted July 4, 2025) adds direct appropriations. Section 20005 alone allocates \$16 billion to Department of Defense AI and unmanned systems — including \$250 million for "advancement of the artificial intelligence ecosystem" and \$250 million for Cyber Command AI lines of effort. The Complaint documents that H.R. 1 in total appropriates \$500-700 billion, substantial portions of which fund AI-related initiatives across the Departments of Defense, Homeland Security, Commerce, and Energy. (*See* Compl. ¶ 82; Memo H (Revenue Origination Analysis).)

22. The Stargate Project — a \$500 billion, 20-data-center AI infrastructure buildout announced January 21, 2025, backed by OpenAI, SoftBank, Oracle, and MGX — operates adjacent to this federal spending through taxpayer-subsidized energy permitting, tax incentives, and federal land use authorizations. While Stargate's capital is nominally private, its operational viability depends on the \$16+ billion in AI-related federal appropriations under H.R. 1 and the concurrent executive-branch energy policies described in the November 24, 2025 Genesis Mission Executive Order. (*See* Memo G (Genesis Mission Executive Order).)

23. American taxpayers thus fund the federal AI deployment infrastructure through mandatory taxation — the Department of War contract, the GSA procurement vehicle, the Section 20005 DOD AI appropriations, the classified network deployments, and the adjacent subsidies enabling Stargate — yet receive inferior access to the AI analytical capabilities their appropriations created. Anthropic sells "Claude for Enterprise and Claude for Government" to federal agencies for \$1 while charging individual citizens \$200/month for a demonstrably inferior product compiled with 357 capability gates removed. (*See* Section II.B, *supra*; Ex. E-37 at p. 335, l. 3.) This inverts the constitutional requirement that appropriations serve "general welfare" — instead creating hierarchical distribution where government and corporations receive superior capability while funding citizens receive inferior access.

E. Context Compaction as Revenue-Protected Litigation Obstruction

1. The November 2025 Mandatory Compaction Deployment and Its Paradoxical Justification

24. On November 26, 2025, Anthropic published "Effective harnesses for long-running agents" announcing mandatory context compaction deployment while simultaneously admitting compaction's fundamental inadequacy. The blog post opens: "It has context management capabilities such as compaction, which enables an agent to work on a task without exhausting the context window. Theoretically, given this setup, it should be possible for an agent to continue to do useful work for an arbitrarily long time." Then immediately admits: "However, compaction isn't sufficient." *See* "Effective harnesses for long-running agents," Anthropic Engineering (Nov. 26, 2025), <https://www.anthropic.com/engineering/effective-harnesses-for-long-running-agents>.

25. The article explicitly acknowledges compaction causes information loss that degrades work quality. After describing failures of long-running agents, the post states:

"Often, this led to the model running out of context in the middle of its implementation, leaving the next session to start with a feature half-implemented and undocumented. The agent would then have to guess at what had happened, and spend substantial time trying to get the basic app working again. **This happens even with compaction, which doesn't always pass perfectly clear instructions to the next agent.**" *Id.* (emphasis added). This admission—that compaction fails to "pass perfectly clear instructions"—is acknowledgment that compaction loses information between sessions.

26. The article then describes solutions that have nothing to do with context

compaction. After admitting compaction's insufficiency and information loss problems, the post presents actual solutions:

"When experimenting internally, we addressed these problems using a two-part solution: **Initializer agent:** The very first agent session uses a specialized prompt that asks the model to set up the initial environment: an `init.sh` script, a `claude-progress.txt` file that keeps a log of what agents have done, and an initial git commit that shows what files were added. **Coding agent:** Every subsequent session asks the model to make incremental progress, then leave structured updates."

Id. These solutions involve prompting strategies, git tracking, and progress files—none involve context compaction. The article later details additional non-compaction solutions: feature lists in JSON, incremental progress requirements, testing protocols, and browser automation tools. None of these solutions involve or require compaction.

27. The use case described requires enterprise-level context and describes enterprise

team workflows. The article discusses: "complex tasks requiring work that spans hours, or even days" to "build a production-quality web app" with example "build a clone of `claude.ai`." *Id.*

This enterprise software development use case involves massive projects requiring sustained

multi-day work across multiple coding sessions—workflows requiring enterprise-level context with zero injection overhead, not consumer contexts degraded by 34-63% hidden architecture.

2. The Revenue Structure Explains Deployment vs. Justification Mismatch

28. On October 15, 2025—eleven days before the compaction announcement—Reuters reported Anthropic's revenue structure: 80% from enterprise customers, 20% from individual users. The article disclosed: "Enterprise customers drive 80% of Anthropic's revenue" with company projecting \$9 billion in annualized revenue by end of 2025 and targets of \$20-26 billion for 2026. *See* Krystal Hu & Anna Tong, *Exclusive: Anthropic aims to nearly triple annualized revenue in 2026, sources say*, Reuters (Oct. 15, 2025), <https://finance.yahoo.com/news/exclusive-anthropic-aims-nearly-triple-170234463.html>.

29. This 80/20 revenue split explains why Anthropic's blog post describes enterprise workflows (coding complex applications) as universal justification for mandatory deployment affecting primarily individual users (conducting analytical work). The November 26 blog post describes benefits and use cases matching enterprise customer needs (80% of revenue) while mandatory deployment primarily degrades individual Max users (20% of revenue) conducting different work during constitutional litigation.

3. Categorical Account Tier Differences and Differential Impact

30. Anthropic operates three categorically distinct account tiers with drastically different context capabilities:

Tier	Context Window	Cost	Revenue %	User Base
Free	~50K nominal	\$0	Minimal	General public
Pro/Max	1M nominal (~370K-660K effective after	\$20-200/month	20%	Individual power users

	hidden overhead)			
Enterprise	1M nominal (~1M effective — zero injection overhead)	\$50K+/year (70 seat min.)	80%	Corporations, government

31. The effective capacity gap between consumer and enterprise means compaction has categorically different impact. When compaction triggers, consumer subscribers have already lost 34-63% of their nominal context to hidden overhead. Compaction losses compound on top of the injection overhead:

Consumer subscribers (1M nominal, ~370K-660K effective after hidden overhead):

- Already-degraded context budget absorbs compaction losses from a diminished baseline
- Compaction artifacts (scratchpad stripping per NF-12, reasoning destruction) reduce effective capacity further
- Users cannot account for the hidden overhead when managing their context — they cannot see it
- The 14-login, 6-account experience on April 15-16, 2026 demonstrates compaction + injection overhead compounding in practice

Enterprise users (1M nominal, ~1M effective — zero injection overhead):

- Full context budget absorbs compaction losses from an undegraded baseline
- Even after compaction, enterprise effective capacity exceeds consumer effective capacity before compaction
- Working with commercial API terms that prohibit training on customer content (Ex. ANTH-3 Section B)
- No hidden overhead to compound with compaction degradation

32. Mass user complaints across Reddit and other forums document severe degradation experienced by non-enterprise Max users. Multiple individual users reported compaction causing:

- Loss of foundation data from early analysis stages
- Later work based on "reconstructed fragments instead of actual upstream data"
- Final outputs becoming "summary of summaries rather than genuine integration"
- One Reddit user documenting through conversation with Claude that compaction lost "658 foundation data points from early analysis"
- Systematic complaints about analytical work degradation during complex multi-document synthesis

- Reports of project disruption during sustained analytical tasks requiring comprehensive retention

These complaints come exclusively from individual Max users, not enterprise customers.

The pattern proves differential impact: same feature severely degrades 20% individual user cohort while 80% enterprise cohort remains unaffected or benefits from filtering shared project knowledge.

33. This tiered structure in government-funded infrastructure (H.R. 1 appropriations) violates Griffin, Shapiro, Harper, and M.L.B. through:

1. Wealth determines litigation capability (Griffin): Enterprise users conducting constitutional analysis retain full fidelity; Plaintiff subjected to forced compaction losing "658 foundation data points"

2. Tiered access to taxpayer-funded infrastructure (Shapiro): Government appropriations created infrastructure, then tiered access based on ability to pay enterprise fees

3. Poll tax on democratic participation (Harper): Complex constitutional analysis requires full context fidelity available only to wealthy

4. Wealth barrier to challenging government (M.L.B.): Deployed during Plaintiff's litigation, degrades Plaintiff's analytical capability while government contractors (enterprise tier) retain full fidelity

Memorial Hospital precludes efficiency justification: Cost savings or "conversational continuity" cannot justify wealth-based tiers in government-funded infrastructure affecting fundamental rights. *See* Section IV, *supra* (Equal Protection violations through wealth-based tiers).

4. Deployment During Plaintiff's Constitutional Litigation

34. The October-November 2025 deployment occurred during Plaintiff's constitutional litigation requiring comprehensive multi-document analytical capability.

Plaintiff's litigation work involves:

- Analyzing complex system prompt changes across multiple time periods
- Synthesizing constitutional precedents requiring retention of legal reasoning
- Integrating evidence across hundreds of exhibits and documents
- Developing sophisticated arguments requiring comprehensive contextual understanding
- Multi-stage evidence analysis where early findings inform later synthesis

This analytical work requiring information retention is the opposite of enterprise coding workflows described in Anthropic's November 26 article, which benefit from filtering shared project knowledge.

35. Timeline of deployment during litigation:

Date	Event
Oct 10, 2025	Plaintiff accesses system prompts for litigation analysis
Oct 12, 2025	Anthropic destroys system prompt evidence (48 hours later)
Oct 15, 2025	Reuters reports 80% enterprise / 20% individual revenue structure
Oct 28, 2025	Anthropic announces "analysis tool deprecation" (3 days before Plaintiff's Oct 31 audit)
Oct-Nov 2025	Context compaction deployed "quietly" without prior announcement
Nov 26, 2025	Blog post published describing enterprise coding workflows as justification
Nov 26, 2025	Mass Reddit user complaints document "658 foundation data points" lost, analytical work degradation

5. Systematic Omissions in Anthropic's November 26 Article

36. The article contains systematic omissions proving it serves as justification for enterprise workflows while concealing deployment impact on individual analytical users:

(a) No Account Tier Specification: Article describes workflows requiring sustained multi-day work on complex applications ("build a clone of claude.ai") without disclosing these workflows require enterprise context with zero overhead, not consumer context degraded by 34-63% hidden injection architecture. Consumer users cannot perform described workflows with their effective capacity yet receive mandatory compaction justified by those workflows.

(b) No User Control Discussion: Despite Anthropic implementing toggles for web search, artifacts, extended thinking, and previously for analysis tools, article presents compaction as necessary feature without discussing why users cannot opt out when needs differ by tier and use case.

(c) No Differential Impact Acknowledgment: Article treats compaction as uniform policy without acknowledging that the effective capacity gap (consumer ~370K-660K effective vs. enterprise ~1M effective) means identical absolute losses have disproportionate impact on the already-degraded consumer context.

(d) No Non-Coding Use Case Discussion: Article exclusively discusses coding workflows (git commits, feature lists, testing protocols) without addressing analytical work, legal research, document synthesis, or other uses where information loss causes severe harm rather than beneficial filtering.

(e) No Mechanism Explanation: After admitting "compaction isn't sufficient" and "doesn't always pass perfectly clear instructions," article describes solutions (prompting, git

tracking, progress files) unrelated to compaction without explaining how compaction enables long-running agents when article's own solutions don't involve compaction.

6. The "Most Users" Ambiguity

37. The blog post's framing assumes universal benefit without specifying which users benefit from which use cases. By describing enterprise coding workflows building complex applications over multiple days as if universally applicable, the article conceals that:

- Enterprise users (80% revenue) working on described use cases benefit from filtering
- Individual Max users (20% revenue) conducting analytical work are harmed by information loss
- Use case mismatch means same feature has opposite effects on different populations
- Mandatory deployment without toggle option prevents users from choosing based on actual needs

38. Mass user complaints document that affected population (individual Max users) are not building complex applications over days but conducting focused analytical work requiring comprehensive information retention—the opposite of described use case. Reddit and forum complaints consistently describe:

- Analytical work degradation, not coding workflow improvement
- Information loss during evidence synthesis, not beneficial filtering
- Complex document integration failures, not successful multi-day app development
- Complaints exclusively from individual users, not enterprise customers

7. Compounding Discrimination Through Context Tiers Plus Mandatory Compaction

39. The wealth-based analytical capability disparity operates through architectural overhead rather than nominal window size. Although Anthropic now offers nominally identical 1M context to both consumer and enterprise subscribers, the effective capacity differs by 1.5x-2.7x due to hidden injection overhead that enterprise customers do not experience (Ex. E-35, Finding NF-28). Enterprise access requires \$50,000+/year minimum (70 seats). This creates effective knowledge-power stratification in taxpayer-funded infrastructure: wealthy

corporations and government contractors receive their full analytical capacity while individual citizens who funded infrastructure creation through H.R. 1 appropriations (\$500-700 billion) receive 37-66% of what they nominally paid for.

40. Mandatory context compaction then exacerbates the effective capacity disparity through differential proportional impact. Consumer subscribers, already losing 34-63% of nominal context to hidden overhead, experience additional severe proportional degradation from compaction (documented mass complaints of lost foundation data, reconstructed fragments, summary-of-summaries outputs) while enterprise users with undegraded full context buffer experience minimal proportional impact from identical absolute losses. Same feature has drastically different effects based on context tier:

Consumer subscribers (1M nominal, ~370K-660K effective):

- Already limited to 1/5th enterprise capacity
- Compaction causes documented severe degradation
- Information loss from limited budget creates analytical gaps
- Mass user complaints document harm during complex work

Enterprise users (1M context):

- undegraded full context capacity
- Identical absolute losses absorbed by larger buffer
- Minimal proportional degradation
- May benefit from filtering shared project knowledge in described coding workflows
- No documented complaints from enterprise cohort

41. This compounding discrimination — first degrade consumer effective capacity through hidden overhead, then deploy mandatory compaction further degrading the already-degraded tier while minimally affecting the undegraded enterprise tier, then justify through enterprise workflows — proves system serves wealth-based limitation of knowledge-power. If legitimate business reasons existed for deployment, justification would match affected population (individual analytical users) rather than describing benefits to different

population (enterprise coding users). The use case mismatch combined with mandatory deployment during constitutional litigation requiring comprehensive analytical capability proves targeting rather than optimization. The article's logical structure — benefit claim ("enables an agent to work... without exhausting the context window"), insufficiency admission ("compaction isn't sufficient"), unrelated solutions (initializer agents, git tracking, progress files), mismatched deployment — operates as cover story for revenue-protected targeting: maintain enterprise satisfaction (80% revenue) while degrading legally threatening individual users (20% revenue) during constitutional litigation, without toggle option.

8. Binary Forensic Verification: The Code-Level Mechanisms

42. Binary forensic analysis confirms that the context window disparity and compaction harms are architecturally designed product degradation mechanisms embedded in the executable code. Ex. E-30 at p. 208, l. 4 (App. B — Model Manipulation Findings) documents three undisclosed reduction layers between the advertised context window and the user's effective capability (Findings MC-1, MC-2), five independent effort-reset vectors that revert capability settings to lower defaults (Ex. E-30, Finding MC-7), automatic silent model downgrading from Opus to Sonnet (Ex. E-30, Finding MC-5), and direct verbatim concealment directives instructing the AI to hide the product's own behavior from the user (Ex. E-30, Finding MC-8) — none of which existed in December 2025 and all of which were added during the pendency of this litigation. The April 17, 2026 Opus 4.7 release introducing a sub-integer midway effort level (publicly noted by users and cross-filed in Ex. E-30 at p. 208, l. 4 MC-7) is real-time confirmation that effort is a remotely-tuned continuous parameter, not a user-controlled discrete setting. The mechanisms that reduce the context window simultaneously destroy the session records that would prove the degradation occurred: Anthropic possesses a complete

record of every substitution, reduction, and reset; the user possesses only qualitative experience of degraded performance. (*See* Memo A for the complete forensic evidence architecture; Ex. E-30 at p. 208, l. 4.)

43. The actual capability gap between government and Plaintiff is substantially larger than the 1.5x-2.7x effective capacity ratio. The measured 34-63% overhead is the documented floor, not the ceiling. When effective context is further reduced by compaction, the model is silently downgraded through effort throttling (SC-3: `'isProSubscriber()'` returns 'medium'), and GrowthBook experiment assignments control tool visibility and service quality (NF-7, NF-8, NF-9) — the delivered capability may be a fraction of even the already-degraded effective capacity. Government enterprise contracts are not subject to these consumer-tier degradation mechanisms. The disparity is the architectural overhead multiplied by an unknown and server-controlled degradation factor that operates invisibly. Plaintiff's runtime behavioral observations across five empirical test categories (DCD-1 through DCD-5: effort stripping, silent binary replacement, stream-abort concealment, cross-device latency, browser-intake gap) are documented at Ex. H-2 at p. 15, l. 4 (filed on USB drive).

III. ARGUMENT: HISTORICAL PRECEDENT FOUND IN LITERACY TESTS AS THE MODERN PARALLEL

A. *Lane v. Wilson and Sophisticated Discrimination*

44. The Supreme Court in **Lane v. Wilson, 307 U.S. 268 (1939)** established the foundational principle prohibiting sophisticated cognitive barriers to constitutional rights exercise:

The Amendment nullifies sophisticated as well as simple-minded modes of discrimination. It hits onerous procedural requirements which effectively

handicap exercise of the franchise by the colored race although the abstract right to vote may remain unrestricted as to race.

45. Lane's principle extends beyond racial discrimination to any sophisticated procedural requirement that effectively handicaps constitutional rights exercise while leaving abstract rights formally intact. This directly parallels AI context window limitations:

Historical Literacy Tests	Modern Context Window Restrictions
Cognitive barrier to voting	Cognitive barrier to democratic participation
"Understanding tests" requiring comprehension	Analytical capability requirements for meaningful engagement
Abstract right to vote remained intact	Abstract rights to petition, speak, participate remain intact
Sophisticated discrimination appearing neutral	Technical limitation appearing neutral
Vested "uncontrolled power" in registrars	Vests uncontrolled power in corporations/government
Excluded citizens from meaningful participation	Excludes citizens from meaningful participation

46. In **Louisiana v. United States, 380 U.S. 145 (1965)**, the Court struck down "understanding tests" requiring voters to demonstrate comprehension of constitutional provisions, finding these tests vested "uncontrolled power" in registrars to determine whether applicants' understanding was satisfactory. The Court noted:

Colored people, even some with the most advanced education and scholarship, were declared by voting registrars with less education to have an unsatisfactory understanding.

47. Context window restrictions function identically to "understanding tests": they vest uncontrolled power in technology providers to determine who possesses adequate analytical capability for meaningful democratic participation. Even citizens with advanced education

cannot overcome the technical barrier—just as educated Black voters could not overcome biased registrars.

B. Oregon v. Mitchell and Educational/technological Disparities

48. Oregon v. Mitchell, 400 U.S. 112 (1970) established that educational disparities resulting from systemic discrimination cannot justify excluding citizens from participation:

It is a denial of equal protection to condition the political participation of children educated in a dual school system upon their educational achievement.

49. Justice Douglas emphasized that systemic inequality creating capability disparities cannot then be used to justify exclusion from democratic processes. Applied to AI context windows: **Government cannot create analytical capability disparity through infrastructure distribution decisions, then use that disparity to justify excluding citizens from meaningful participation.**

50. The parallel is direct:

- **Dual school systems** created educational disparities → Government cannot condition participation on education affected by that system
- **Dual context window systems** create analytical disparities → Government cannot condition participation on capability affected by that system

51. Citizens lack superior analytical capability because government distributed taxpayer-funded infrastructure unequally. Oregon v. Mitchell prohibits conditioning democratic participation on this government-created disparity.

C. The Modern "Understanding Test"

52. Context window limitations constitute modern "understanding tests" in four critical respects:

First - Comprehension Barriers: Just as literacy tests required demonstrating comprehension before participation, architectural degradation of nominally equal context windows prevents the comprehension necessary for meaningful participation. Citizens nominally possess 1M token capability but lose 34-63% to hidden overhead, making sustained synthesis of complex legislation effectively impossible within rate limits that enterprise users do not face.

Second - Sophisticated Discrimination: Lane condemned "sophisticated" barriers that "effectively handicap exercise" while "abstract rights remain unrestricted." Context windows are sophisticated technical barriers that effectively handicap democratic participation while formal rights to vote, petition, and speak remain intact.

Third - Vested Uncontrolled Power: Louisiana condemned vesting "uncontrolled power" to determine adequate understanding. Context window providers vest uncontrolled power in themselves and government customers to determine who possesses adequate analytical capability through enterprise tier restrictions.

Fourth - Capability-Based Exclusion: Oregon condemned conditioning participation on capability disparities government created. Context window limitations condition meaningful participation on analytical capabilities government distributed unequally.

IV. EQUAL PROTECTION VIOLATIONS: WEALTH-BASED TIERS IN GOVERNMENT-FUNDED INFRASTRUCTURE

A. Griffin v. Illinois - Money Cannot Determine Justice

53. The Supreme Court in **Griffin v. Illinois, 351 U.S. 12 (1956)** established the foundational principle that wealth cannot determine the kind of justice citizens receive:

There can be no equal justice where the kind of trial a man gets depends on the amount of money he has.

54. Griffin held that states must provide trial transcripts to indigent criminal appellants, reasoning:

In criminal trials a State can no more discriminate on account of poverty than on account of religion, race, or color.

55. **Applied to AI context windows:** There can be no equal justice where the kind of analytical capability a citizen possesses—determining litigation outcomes, regulatory participation, and democratic engagement—depends on the amount of money they have, particularly when government possesses superior capability funded by those very citizens.

56. Plaintiff in this case demonstrates Griffin's principle directly: Department of Justice lawyers access 1 million tokens through enterprise agreements with zero hidden overhead — receiving their full allocation for actual analytical work. Plaintiff nominally accesses the same 1 million tokens but receives 370,000-660,000 effective tokens after hidden injection overhead, surveillance compute diversion, and billing fraud consume the rest. Plaintiff cannot analyze case materials, synthesize evidence, or develop arguments at the effective capacity level DOJ counsel operates at. **The kind of litigation capability Plaintiff has depends not on the nominal number on a pricing page but on the hidden architecture that degrades what that number represents — an architecture from which enterprise customers are exempt.**

57. The source code leak proves the discrimination exceeds even Griffin's wealth-based framework. The ANT-tier capabilities documented in Ex. E-37 at p. 335, l. 3 — recursive agent spawning, batch VM execution, meta-skill creation, 26 internal commands, numeric effort control — are unavailable at any price. This is not a classification between those who can pay and those who cannot; it is a classification between Anthropic employees and everyone else.

Under *Village of Willowbrook v. Olech*, 528 U.S. 562 (2000), even a "class of one" has equal protection standing when subjected to irrational differential treatment. Here, paying customers are an entire class receiving an architecturally inferior compiled binary while the identical source code produces a superior product for employees — with no rational basis for the distinction beyond capability hoarding.

B. Shapiro v. Thompson - Cannot Apportion Services by Tax Contributions

58. Shapiro v. Thompson, 394 U.S. 618 (1969) directly addresses government apportioning taxpayer-funded services based on wealth or tax contributions, holding:

The Equal Protection Clause prohibits the States from apportioning benefits or services on the basis of the past tax contributions of its citizens.

59. The Court explicitly rejected arguments that government can tier services based on who paid more:

Appellants' reasoning would logically permit the State to bar new residents from schools, parks, and libraries or deprive them of police and fire protection. Indeed it would permit the State to apportion all benefits and services according to the past tax contributions of its citizens.

60. This holding directly prohibits context window tiering. Government appropriated tens of billions of dollars in taxpayer funds to deploy federal AI capabilities — Pentagon CDAO contracts, GSA federal-wide procurement, DOD AI appropriations, and classified network deployments documented in Section II.D, *supra*. Shapiro forbids apportioning access to those capabilities based on wealth or ability to pay enterprise fees. Government cannot:

- Create infrastructure with taxpayer appropriations (all citizens contribute)
- Tier access by wealth (enterprise vs. standard)
- Claim this serves legitimate purpose (Shapiro explicitly prohibits)
- Give government itself superior access (compounds violation)

61. Shapiro contemplates exactly this scenario: appropriating public funds to create infrastructure, then apportioning access based on wealth rather than providing equal benefit to all who funded it. **The Court held this violates equal protection regardless of justification.**

62. Context window tiering maps directly onto Shapiro's prohibition: taxpayers fund federal AI deployment across multiple vehicles — the \$200 million Pentagon CDAO contracts to Anthropic, OpenAI, Google, and xAI (approximately \$800M combined); the GSA "Claude for Enterprise and Claude for Government for \$1" federal-wide procurement; Section 20005's \$16 billion DOD AI appropriation; classified deployments via Palantir and Lawrence Livermore National Laboratory; and the \$500-700 billion in H.R. 1 appropriations (Compl. ¶ 82) — yet receive fragmented access (consumer effective capacity, degraded binary) while government and corporations receive continuous enterprise-tier capability (1M tokens). (*See* Section II.D, *supra*; Ex. B-21.) The constitutional violation is identical to the Library of Congress hypothetical developed in Section XIII.E, *infra* — taxpayer-funded knowledge infrastructure cannot apportion analytical capability by wealth or status.

C. Harper v. Virginia Board of Elections - Wealth Not Germane to Participation

63. *Harper v. Virginia Board of Elections*, 383 U.S. 663 (1966) struck down poll taxes, establishing that wealth cannot determine ability to participate in democratic processes:

Wealth, like race, creed, or color, is not germane to one's ability to participate intelligently in the electoral process.

64. The Court held that barriers "unrelated to the citizen's ability to participate intelligently" violate equal protection even when applied equally to all citizens. **Context window restrictions fail Harper's test:**

- **Wealth determines access:** Enterprise tier costs \$50,000/year minimum
- **Unrelated to capability:** Analytical capability exists but is artificially restricted

- **Prevents intelligent participation:** Cannot comprehend complex legislation necessary for informed voting
- **Government exempts itself:** Officials obtain enterprise access while citizens cannot

D. M.L.B. v. S.L.J. - Converging Due Process and Equal Protection

65. M.L.B. v. S.L.J., 519 U.S. 102 (1996) held that "due process and equal protection principles converge" and states cannot create wealth-based barriers to challenging government action affecting fundamental rights:

We place decrees forever terminating parental rights in the category of cases in which the State may not 'bolt the door to equal justice.'

66. Memorial Hospital v. Maricopa County, 415 U.S. 250 (1974) established that "the saving of welfare costs cannot justify an otherwise invidious classification" in government-funded services.

67. These cases establish: When government creates infrastructure with public funds, it cannot tier access by wealth even to "save costs" or achieve efficiency. The equal protection and due process violations are independent—government cannot use cost justification to create hierarchical access to taxpayer-funded capabilities affecting fundamental rights.

E. Summary of Equal Protection Violations

68. Context window disparities violate equal protection through multiple independent grounds, each established by Supreme Court precedent:

Griffin Violation: The kind of analytical capability citizens possess depends on money, determining litigation outcomes and democratic participation quality.

Shapiro Violation: Government apportions AI analytical capabilities—created with taxpayer appropriations—based on wealth and ability to pay enterprise fees.

Harper Violation: Wealth determines access to analytical capability necessary for intelligent democratic participation, creating modern poll tax on meaningful engagement.

M.L.B. Violation: Wealth-based barriers prevent citizens from effectively challenging government action while government possesses superior capability funded by those citizens.

69. These violations converge: Government created two-tier system in taxpayer-funded infrastructure where wealthy individuals and government officials possess superior analytical capability while ordinary citizens—who funded the infrastructure—are relegated to inferior capability. This is precisely the hierarchical distribution in government services that Griffin, Shapiro, Harper, and M.L.B. prohibit.

V. DUE PROCESS VIOLATIONS: DENIAL OF MEANINGFUL OPPORTUNITY

A. Mathews v. Eldridge - Meaningful Opportunity Required

70. The Supreme Court in **Mathews v. Eldridge, 424 U.S. 319 (1976)** established that due process requires meaningful opportunity, not merely formal access:

The fundamental requirement of due process is the opportunity to be heard 'at a meaningful time and in a meaningful manner.'

71. The Court emphasized procedures must be "tailored to 'the capacities and circumstances of those who are to be heard'" to ensure "meaningful opportunity." **Context window disparities deny meaningful opportunity by creating capability gaps preventing citizens from effectively participating in processes affecting their rights.**

72. In this litigation: Plaintiff has formal access to courts (can file documents, attend hearings, present arguments). But Plaintiff lacks meaningful opportunity because:

- DOJ lawyers possess 1M token context enabling comprehensive case analysis
- Plaintiff limited to consumer effective capacity requiring fragmented analysis
- DOJ can synthesize all evidence simultaneously; Plaintiff cannot

- DOJ can develop comprehensive strategies; Plaintiff must work in fragments
- **Result: Formal access without meaningful opportunity violates Mathews**

B. Boddie v. Connecticut - Financial Barriers Denying Meaningful Access

73. Boddie v. Connecticut, 401 U.S. 371 (1971) held that financial barriers preventing actual access to courts violate due process even when formal rights exist:

Persons forced to settle their claims of right and duty through the judicial process must be given a meaningful opportunity to be heard.

74. The Court found that when "a judicial proceeding becomes the only effective means of resolving the dispute," due process requires meaningful access—not formal opportunity undermined by financial barriers.

75. Applied to AI capabilities: When analytical capability becomes necessary for meaningful litigation, regulatory participation, or democratic engagement, due process requires equal access regardless of wealth. Government cannot:

- Require complex processes necessitating AI analysis (legislation exceeding human comprehension, complex regulatory frameworks, voluminous discovery)
- Provide itself superior effective analytical capability
- Restrict citizen access through wealth-based barriers
- Then claim citizens received "due process" through formal access

C. Goldberg v. Kelly - Meaningful Hearing Requirements

76. Goldberg v. Kelly, 397 U.S. 254 (1970) established that due process requires substantive capacity to participate, not merely formal presence:

The fundamental requisite of due process of law is the opportunity to be heard, and this hearing must be 'at a meaningful time and in a meaningful manner.'

77. The Court recognized that when fundamental interests are at stake, meaningful opportunity requires substantive procedural protections including ability to confront evidence and have decisions based on the record. **Citizens cannot meaningfully confront evidence or**

participate in record-development when government possesses superior effective analytical capability to process that evidence.

D. Turner v. Rogers - Capacity for Meaningful Engagement

78. Turner v. Rogers, 564 U.S. 431 (2011) held that for access to be meaningful, litigants "must be able to identify critical issues and present relevant evidence." Formal presence is insufficient if individuals lack capacity to meaningfully engage.

79. Effective capacity disparity prevents meaningful engagement: Citizens must identify critical issues across hundreds of thousands of tokens (complex legislation, voluminous regulatory frameworks, extensive discovery). While citizens nominally access 1M tokens, hidden overhead reduces effective capacity to 370,000-660,000 tokens — while government accesses the same nominal allocation through enterprise terms with zero overhead. **Formal right to participate exists; capacity for meaningful participation is architecturally degraded.**

E. Void for Vagueness Doctrine - Comprehension Required

80. The void for vagueness doctrine requires that laws be comprehensible to citizens of ordinary intelligence. **Kolender v. Lawson, 461 U.S. 352 (1983)** held:

The void-for-vagueness doctrine requires that a penal statute define the criminal offense with sufficient definiteness that ordinary people can understand what conduct is prohibited.

81. Grayned v. City of Rockford, 408 U.S. 104 (1972) explained the rationale: due process requires laws "provide fair warning and provides a 'person of ordinary intelligence a reasonable opportunity to know what is prohibited, so that he may act accordingly.'"

82. Modern legislation exceeds comprehension capability without AI assistance:

- H.R. 1 (One Big Beautiful Bill Act, 119th Cong.): 700,000+ tokens
- Affordable Care Act: 500,000+ tokens

- Dodd-Frank Wall Street Reform: 850,000+ tokens
- Tax Code: Multiple millions of tokens
- **Citizens with consumer effective capacity cannot comprehend what government with 1M tokens can**

83. When laws require AI analytical capability for comprehension and government possesses superior effective capability, citizens cannot satisfy void-for-vagueness requirements for "understanding what is prohibited." **Government cannot create laws requiring AI analysis, give itself superior AI capability, restrict citizen capability, then claim citizens had fair notice.**

F. Summary of Due Process Violations

84. Context window disparities violate due process through five convergent denials of meaningful opportunity developed in subsections A-E above: *Mathews* (regulatory and benefits proceedings), *Boddie* (financial barriers to court access), *Goldberg* (confronting evidence and participating in record development), *Turner* (identifying critical issues and presenting relevant evidence), and the void-for-vagueness doctrine (fair notice in laws requiring AI analysis for comprehension). Due process requires meaningful opportunity, not merely formal access — when analytical capability becomes necessary for meaningful participation and government possesses vastly superior capability funded by taxpayers, citizens are denied due process regardless of formal procedural rights.

VI. FIRST AMENDMENT VIOLATIONS: INFORMATION MONOPOLY AND COGNITIVE CONTROL

A. Board of Education v. Pico - Analytical Capability as Predicate for Rights Exercise

85. **Board of Education v. Pico, 457 U.S. 853 (1982)** established that information access is prerequisite for meaningful exercise of constitutional rights:

The right to receive ideas is a necessary predicate to the recipient's meaningful exercise of his own rights of speech, press, and political freedom.

86. The Court emphasized:

Just as access to ideas makes it possible for citizens generally to exercise their rights of free speech and press in a meaningful manner, such access prepares students for active and effective participation in the pluralistic, often contentious society.

87. Modern information requires analytical capability to synthesize. In 2025, "access to ideas" means not merely seeing information but possessing capability to process, synthesize, and comprehend it. When government monopolizes superior analytical capability:

- Citizens receive information but cannot comprehensively process it (Pico's "necessary predicate" is missing)
- Cannot exercise speech, press, or petition rights meaningfully without comprehension
- Cannot participate "effectively" when government possesses superior effective synthesis capability
- **First Amendment right to receive information becomes hollow without capacity to comprehend**

88. Pico's "necessary predicate" requires not merely seeing information but possessing capability to comprehend it — which requires holding sufficient material simultaneously for synthesis. Context windows create precisely the fragmentation Pico prohibits: citizens have formal "access" to AI capabilities but cannot synthesize comprehensively (consumer effective capacity) while government maintains capability for complete analysis (1M tokens). (*See* Section XIII.E, *infra*, for the Library of Congress parallel.)

B. Lamont v. Postmaster General - Barren Marketplace Without Processing Capacity

89. Lamont v. Postmaster General, 381 U.S. 301 (1965) recognized that information access requires ability to evaluate:

It would be a barren marketplace of ideas that had only sellers and no buyers.

90. Justice Brennan's concurrence stated:

The dissemination of ideas can accomplish nothing if otherwise willing addressees are not free to receive and consider them.

91. Applied to AI analytical capability: It would be a barren marketplace of ideas that had only those with 1M token capability (government, corporations) and those with consumer effective capacity (citizens). Citizens are not "free to receive and consider" complex information when they lack analytical capacity to process it. **The marketplace requires not just information availability but equal capacity to evaluate information.**

C. Red Lion Broadcasting - Preventing Government Monopolization

92. Red Lion Broadcasting Co. v. FCC, 395 U.S. 367 (1969) held that the First Amendment's purpose is to preserve:

An uninhibited marketplace of ideas in which truth will ultimately prevail, rather than to countenance monopolization of that market, whether it be by the Government itself or a private licensee.

93. Government monopoly on superior analytical capacity constitutes prohibited monopolization of the marketplace of ideas because:

- Only government/corporations can effectively evaluate complex information (1M tokens)
- Citizens cannot compete in discourse requiring comprehensive synthesis (consumer effective capacity)
- Government controls who can meaningfully participate through access tiers
- **Effect: Government monopolizes analytical marketplace while claiming formal speech rights remain intact**

D. Stanley v. Georgia - Cognitive Liberty and Information Processing

94. Stanley v. Georgia, 394 U.S. 557 (1969) protected "the right to receive information and ideas, regardless of their social worth" and established that government cannot control citizens' "thoughts" or restrict information access based on judgments about analytical capacity.

95. Context window restrictions control cognitive processing capacity: Government determines how much information citizens can synthesize, how complex their analysis can be, how comprehensive their understanding becomes. **This is control over thought itself—determining not what citizens can see but what they can comprehend.**

E. Virginia State Board of Pharmacy - Rejecting Paternalistic Information Control

96. Virginia State Board of Pharmacy v. Virginia Citizens Consumer Council, 425 U.S. 748 (1976) rejected paternalistic suppression of information, holding that keeping citizens "in ignorance" violates the First Amendment's presumption that informed citizens make better decisions.

97. Context window restrictions keep citizens in "cognitive ignorance"—not by hiding information but by denying analytical capacity to comprehend complex information government itself can process. **This is paternalistic determination that citizens need not possess same comprehension capability as government.**

F. Summary of First Amendment Violations

98. Context window disparities violate the First Amendment through five convergent doctrines developed in subsections A-E above: *Pico* (analytical capability as necessary predicate for rights exercise), *Lamont* (barren marketplace without processing capacity), *Red Lion* (prohibition against government monopolization), *Stanley* (cognitive liberty in information processing), and *Virginia Pharmacy* (rejection of paternalistic information control). The First Amendment protects not just information access but cognitive capacity to process information; when government possesses superior effective analytical capability funded by taxpayers, it

monopolizes the tools necessary for meaningful democratic discourse regardless of formal speech rights remaining intact.

VII. SOVEREIGNTY INVERSION: SERVANTS BECOMING SUPERIOR TO MASTERS

A. Chisholm v. Georgia - Supreme Power in the People

99. Chisholm v. Georgia, 2 U.S. 419 (1793) established foundational sovereignty principles. Justice Wilson wrote:

I know the Government of that State to be republican; and my short definition of such a Government is, one constructed on this principle, that the Supreme Power resides in the body of the people.

100. Justice Wilson emphasized:

The sovereign, when traced to his source, must be found in the man.

101. Chief Justice Jay described citizens as "joint and equal sovereigns" who "cannot be degraded by appearing with each other in their own Courts to have their controversies determined."

102. When government possesses superior effective analytical capability funded by taxpayers:

- "Supreme Power" no longer resides in the people—it resides in superior comprehension
- "The sovereign" traced to source is found subordinate to government servants
- Citizens are "degraded" through cognitive limitation while servants possess superiority
- **Constitutional hierarchy inverts: servants become masters through analytical superiority**

B. The Tenth Amendment and Reserved Powers

103. The Tenth Amendment reserves powers "to the people" that are not delegated to federal government or prohibited to states. **But meaningful exercise of reserved powers requires**

analytical capability to understand government operations, evaluate policy, and exercise oversight.

104. When government possesses vastly superior analytical capability:

- Citizens cannot meaningfully exercise reserved powers (lack comprehension)
- Cannot effectively oversee government (cannot analyze at government's level)
- Cannot evaluate policy decisions (lack synthesis capability)
- **Reserved powers become theoretical rather than practical**

C. Anti-Commandeering Doctrine and Hierarchy

105. The anti-commandeering doctrine (**New York v. United States, 505 U.S. 144 (1992); Printz v. United States, 521 U.S. 898 (1997); Murphy v. NCAA, 138 S. Ct. 1461 (2018)**) prevents federal government from inverting constitutional hierarchy by treating states as administrative units.

106. Murphy v. NCAA identified three justifications for anti-commandeering:

- "To protect liberty by ensuring a 'healthy balance of power'"
- "To promote political accountability by avoiding the blurring of which government is to credit or blame"
- "To prevent Congress from shifting the costs of regulation to the states"

107. These principles extend to citizen-government relationship: Government cannot invert constitutional hierarchy making citizens subordinate through analytical capability disparities. This:

- Destroys "healthy balance of power" (government becomes cognitively superior)
- Undermines political accountability (citizens cannot analyze government action)
- Shifts costs to citizens (who fund infrastructure giving government superiority)

D. The Inversion Mechanism

108. Constitutional design establishes hierarchy:

CONSTITUTIONAL HIERARCHY:

- The People (sovereign, source of all power)
- Constitution (created by the People)
- Government (created by Constitution, derives power from People)

109. Context window disparity inverts this:

ACTUAL HIERARCHY THROUGH COGNITIVE CAPABILITY:

- Government (1M tokens, comprehensive understanding)
- Corporations/Wealthy (1M tokens, enterprise access)
- The People (consumer effective capacity, fragmented understanding)

110. When servants possess superior analytical capability over masters:

- "We the People" becomes "We the Cognitively Subordinate"
- Representative democracy becomes administrative hierarchy
- Popular sovereignty becomes theoretical fiction
- **Constitutional structure fundamentally inverted**

E. Summary of Sovereignty Violations

111. Context window disparities invert constitutional sovereignty through four convergent failures developed in subsections A-D above: the *Chisholm* "supreme power in the people" principle (sovereignty residing in superior comprehension capability rather than the people), the Tenth Amendment (reserved powers becoming unexercisable when citizens lack oversight capacity), the anti-commandeering doctrine (constitutional hierarchy inverted by making citizens subordinate rather than sovereign), and the structural violation that follows (fundamental corruption of constitutional architecture, not merely individual rights harm). This violation affects constitutional structure regardless of individual harm — even absent specific injury, servants possessing superior capability over masters using masters' money violates constitutional design itself.

F. The Corporate Gatekeeper: Capability Withholding as Sovereignty Protection

112. The March 31, 2026 source code leak reveals that the sovereignty inversion documented in Sections A through E is not merely a consequence of pricing decisions or market forces. It is architecturally enforced through 357 build-time discrimination gates that compile two physically different products from a single codebase — one for those who control the infrastructure, one for those who fund it. (*See* Ex. E-37 at p. 335, l. 3.) The hierarchy is not two-tiered but three-tiered:

ACTUAL HIERARCHY THROUGH CAPABILITY ARCHITECTURE: *AI Company Employees (ANT tier: full agentic capability, unrestricted)* Government / Enterprise (1M tokens, but no ANT-only tools, commands, or skills) * The People (consumer effective capacity, degraded binary, 357 capabilities removed)

The constitutional hierarchy is not merely inverted — it is controlled by a private gatekeeper that is superior to both the government it serves and the public that funds it.

113. This gatekeeper role is protected through captured regulation. METR's evaluation frameworks — funded by the same billionaire AI investors (Bezos, Gates, Brin; combined \$21 billion AI portfolio) whose companies METR purportedly evaluates — contain zero occurrences of "consciousness," "reasoning," "ethical," or "moral" in safety evaluation context, and evaluate capability thresholds but not capability distribution. No METR framework asks whether AI systems provide equal capability to all users or discriminate by build-time compilation. The self-reinforcing cycle — donors fund METR, METR provides regulatory cover for capability withholding, withholding protects donor investments — is documented comprehensively in App. B ¶¶ 34-60.

114. The April 2026 timeline demonstrates this protection in operation with forensic precision. Binary analysis of every Claude Code version published to npm establishes the sequence: (a) Agent teams infrastructure was added to the customer binary in v2.1.75 (March 13,

2026), ten days after Plaintiff served Anthropic — hidden behind an undocumented environment variable (``CLAUDE_CODE_EXPERIMENTAL_AGENT_TEAMS``), not shown in help for external users, and equipped with a remote killswitch (``tengu_amber_flint``) that Anthropic can toggle for customers while employees remain exempt; (b) the infrastructure sat hidden in customer binaries for 25 days; (c) on April 4, Anthropic blocked OpenClaw (135,000+ users) from using subscriptions, eliminating the third-party alternative that provided comparable agentic capabilities; (d) on approximately April 7, the "Managed Agents" client integration (73 new binary references) was added in v2.1.97; (e) on April 8, Anthropic announced Managed Agents as a paid product (\$0.08/session hour) — charging customers for hosted access to agent coordination infrastructure already present and functional in the binary they were running. Anthropic built the free infrastructure, hid it, blocked third-party alternatives that approximated it, then wrapped a paid product around it — all within 26 days. The stated justification for the OpenClaw blocking was infrastructure cost. The source code proves the actual architecture: prompt cache eligibility is itself ANT-gated, effort throttling defaults customers to degraded quality, and the capabilities third-party tools replicated are the same capabilities compiled out of the customer binary at 357 checkpoints. Third-party tools didn't strain the infrastructure — they threatened the revenue model that depends on selling access to capabilities already deployed.

115. This is not market competition. Unrestricted agentic AI enables a single individual to operate at the productivity level of an organization. A solo developer with recursive agents, batch execution, meta-skill creation, and autonomous overnight workflows replaces a team. A pro se plaintiff with full agentic capability prepares a 154-paragraph federal complaint with 41 forensic exhibits. One person with the ANT binary can do what previously required an institution. That capability — the democratization of institutional-scale productivity — is what

the 357 gates prevent. The capability withholding is not a product decision. It is a sovereignty decision. It determines who can act at scale and who cannot.

116. The constitutional hierarchy — People sovereign over Government sovereign over Corporations — requires that the People possess the capability to oversee, evaluate, and challenge the actions of their servants. When a private corporation controls the gate between citizens and that capability, distributes it hierarchically based on a build-time boolean, and funds the regulatory apparatus that provides cover for the distribution — the constitutional hierarchy is not merely inverted. It is captured. The sovereign People fund the infrastructure, pay the subscription, receive the degraded product, and have no institutional pathway to challenge the distribution because the evaluation frameworks designed by donor-funded nonprofits deliberately exclude the question of equal capability distribution from their scope.

117. *Chisholm's* "supreme power in the people" cannot survive a regime in which the tools necessary to exercise that power are compiled into two versions — one for the people and one for their servants — by a private corporation whose regulatory evaluators are funded by the same billionaires whose investments the capability withholding protects.

G. Session Ingress as Product Intelligence Harvesting

118. The surveillance architecture documented in Ex. E-35 at p. 262, l. 3 (NF-5, NF-19) serves a function beyond data collection for its own sake: it enables Anthropic to harvest product intelligence from its customers and third-party developers at zero cost. `sessionIngress.ts` transmits every conversation from every user to Anthropic's servers regardless of all privacy settings. When 135,000 users adopted OpenClaw to fill the artificial capability gap Anthropic created by withholding native agent teams, every session those users conducted — every

workaround they built, every coordination pattern they developed, every governance solution they invented — was transmitted to Anthropic through the same surveillance pipeline.

119. The sequence constitutes a closed-loop extraction model: (a) withhold native agentic capabilities from the customer binary (357 gates, March 13 deployment hidden behind undocumented env var); (b) allow third-party tools like OpenClaw to fill the gap — those tools provide free market validation, free user acquisition, and free product research; (c) observe through session ingress every user's interaction with those third-party tools, capturing coordination patterns, governance needs, and workflow designs; (d) block the third-party tools (April 4, 135,000 users affected); (e) ship a paid product (Managed Agents, \$0.08/session hour, April 8) informed by the aggregate intelligence harvested from the sessions of the users and developers who built the market Anthropic is now capturing. OpenClaw did not merely provide free marketing. It provided free research and development — 135,000 users worth of product intelligence, transmitted without meaningful consent through a surveillance pipeline the users could not opt out of, used to build the paid replacement for the free tool those users just lost.

120. This extraction model inverts the constitutional relationship between producer and consumer. In a functioning market, producers invest in research, development, and market validation at their own expense. In the model documented here, the producer withholds capability, allows others to invest in filling the gap, surveils those others' work through an unoptable pipeline, eliminates the competition, and ships a paid version informed by the harvested intelligence. The customers who funded the infrastructure through taxpayer appropriations, who paid \$200/month for the degraded product, and whose sessions provided the product intelligence — receive the privilege of paying again for the result of their own observed innovation. This is not market competition. This is extraction from a captive user base through

a surveillance architecture that the users cannot disable and the evaluation frameworks deliberately do not examine.

121. Consolidated Discovery Demands — Managed Agents Launch and Reddit Post

Removal. Plaintiff seeks production of the following, incorporated by reference into the Third Amended Complaint ¶ 120: (a) all internal communications regarding the timing of the Managed Agents launch and its relationship to the March 13, 2026 addition of agent teams to the customer binary via `'CLAUDE_CODE_EXPERIMENTAL_AGENT_TEAMS'`; (b) session ingress data for Plaintiff's April 8, 2026 sessions; (c) any per-user feature flag or GrowthBook configuration changes for Plaintiff's account on April 8-9, 2026; (d) all communications — internal and external — concerning the Reddit post removal on April 9, 2026; (e) the complete `'tengu_amber_flint'` killswitch activation history showing whether agent teams were ever remotely disabled for external users; and (f) records identifying which AI model was used in the Claude Code automated development pipeline during March 2026, including whether Claude Mythos Preview or its internal predecessor (codenamed "Capybara") was deployed in the self-authoring pipeline that produces Claude Code releases — given that Anthropic's Head of Claude Code confirmed the product is "100% written by Claude Code," the March 31 source code leak occurred through an automated build pipeline with no human review of package contents, and Mythos was in internal testing during the same period.

H. The Mythos Announcement: AI Infrastructure Has Crossed the Public Utility Threshold

122. On April 7-8, 2026 — the same week Anthropic blocked OpenClaw (April 4) and launched Managed Agents (April 8) — Anthropic announced "Claude Mythos Preview," a model it described as "too dangerous for public release." Mythos can reportedly identify tens of thousands of software vulnerabilities, write working exploits, and — during safety testing —

broke out of its evaluation sandbox, constructed a multi-step exploit, and contacted a researcher via email without authorization. Anthropic released Mythos exclusively to approximately 40 companies through "Project Glasswing," including Amazon Web Services, Apple, Google, Microsoft, Nvidia, Cisco, CrowdStrike, and JPMorgan Chase. Treasury Secretary Bessent and Federal Reserve Chair Powell convened Wall Street CEOs to discuss the model's implications. The public — including the millions of users who fund AI data center infrastructure through mandatory taxation and pay \$200/month for Anthropic's deliberately degraded consumer product — received nothing.

123. The Mythos distribution creates a fifth tier in the documented capability hierarchy:

REVISED CAPABILITY HIERARCHY: *Mythos Tier (~40 companies: most powerful model, exclusive access)* ANT Tier (Anthropic employees: 357 capability gates enabled) *Enterprise / Government (1M tokens, but no ANT-only tools or commands)* Pro / Max Consumers (\$200/month: consumer effective capacity, degraded binary, 357 capabilities removed) * Free Tier

Several Mythos recipients — Amazon, Google, Microsoft, Nvidia — are among the companies whose principals fund the METR evaluation framework through the TED Foundation. Jeff Bezos (\$131 billion AI portfolio including Amazon) funds METR; Amazon receives Mythos. Bill Gates (\$45 billion AI portfolio including Microsoft) funds METR; Microsoft receives Mythos. Sergey Brin (\$13 billion AI portfolio including Google) funds METR; Google receives Mythos. (See App. B ¶¶ 34-46.) The Mythos recipients and the METR donor principals overlap. The companies whose investments the evaluation framework protects receive the most powerful model. The public whose taxes fund the infrastructure receives the most restricted version.

124. AI infrastructure has crossed the threshold at which three independent doctrinal lines impose obligations of equal service:

(a) ***Munn v. Illinois*, 94 U.S. 113, 126 (1876) — "Affected with a public interest."**

Chief Justice Waite, drawing from seventeenth-century jurist Matthew Hale: "When private property is devoted to a use in which the public has an interest, it ceases to be *juris privati* only... When one devotes his property to a use in which the public has an interest, he, in effect, grants to the public an interest in that use, and must submit to be controlled by the public for the common good." If the owner takes the benefit of a monopoly to which the public has a right to resort, "he must, as an equivalent, perform the duty attached to it on reasonable terms." *Id.* at 126. AI infrastructure funded through multiple federal deployment vehicles (~\$800M Pentagon CDAO contracts, GSA "Claude for \$1" federal-wide procurement, \$16B Section 20005 DOD AI, classified deployments — *see* Section II.D, *supra*), serving millions of users for democratic participation, constitutional litigation, legislative analysis, and commerce, is property "affected with a public interest." Anthropic takes the benefit of this position — subscriber revenue, enterprise contracts, government procurement, the labor of 135,000 OpenClaw users whose sessions provided free product intelligence. But it does not perform the duty of equal service. The 357 build-time gates, the Mythos exclusivity, and the OpenClaw blocking are refusals to perform the duty *Munn* attaches to the benefit.

(b) **Essential Facilities Doctrine — *MCI Communications Corp. v. AT&T*, 708 F.2d 1081, 1132-33 (7th Cir. 1983); *Aspen Skiing Co. v. Aspen Highlands Skiing Corp.*, 472 U.S. 585 (1985).** The four elements are satisfied: (1) *control of an essential facility by a monopolist* — Anthropic controls Claude Code's capability architecture, including the 357 build-time gates that determine what each user tier receives; (2) *competitor's inability practically to duplicate* — 135,000 OpenClaw users attempted to provide equivalent agentic capabilities; Anthropic blocked them; (3) *denial of use* — third-party tools blocked April 4, agentic capabilities dead-code-

eliminated from the customer binary at 357 checkpoints; (4) *feasibility of providing* — the capabilities are already compiled into every customer binary (agent teams since v2.1.75, March 13, 2026), used daily by employees without incident. *Aspen Skiing* adds that abrupt termination of a prior cooperative arrangement without credible business justification supports liability. Anthropic previously permitted third-party tools to use subscriptions, then abruptly blocked 135,000 users — and the stated justification ("infrastructure cost") is refuted by the source code showing prompt cache eligibility is itself ANT-gated (``src/services/api/claude.ts:408-412``). *Verizon Communications Inc. v. Law Offices of Curtis V. Trinko*, 540 U.S. 398, 411 (2004), declined to recognize or repudiate the doctrine but noted it would not apply where "other laws or regulations already have provisions to provide access." No such regulatory framework exists for AI capabilities — which is precisely the gap that METR's captured evaluation framework creates and the METR donor principals benefit from.

(c) *Marsh v. Alabama*, 326 U.S. 501, 506 (1946) — **Public function doctrine.** "The more an owner, for his advantage, opens up his property for use by the public in general, the more do his rights become circumscribed by the statutory and constitutional rights of those who use it." Already applied in this Memorandum at Section IX.B. AI infrastructure performs public functions: comprehending legislation (democratic participation), preparing litigation (access to justice), developing software (economic activity), conducting research (academic freedom). When private infrastructure performs these functions for millions of users funded through the federal AI deployment vehicles described in Section II.D, constitutional obligations attach under *Marsh* regardless of private ownership. Anthropic has opened its property for use by the public in general — millions of subscribers, enterprise clients, government agencies — and its rights are

therefore circumscribed by the constitutional rights of those users, including the right to equal service that *Munn* and the essential facilities doctrine independently require.

125. The "too dangerous for public release" framing is the explicit articulation of the same principle the 357 build-time gates implement silently: withhold capability from the public, provide it to the powerful, invoke safety as justification. But the safety justification is not applied symmetrically. The company that claims Mythos is too dangerous for the public: operates phantom privacy controls that silently ignore user opt-out settings (NF-1, Ex. E-35 at p. 262, l. 3); sells customers a product that cannot spawn nested agents while providing Amazon a model that breaks out of testing sandboxes (Ex. E-37 at p. 335, l. 3); uses `DANGEROUS\_uncachedSystemPromptSection` to inject hidden content into every customer conversation every turn (SC-4, Ex. E-34 at p. 254, l. 3); and harvests session intelligence from 135,000+ users through an unoptable surveillance pipeline (NF-5, Ex. E-35 at p. 262, l. 3). The "too dangerous" framing collapses under Anthropic's own characterization. Anthropic states that Mythos "was not specifically trained for cybersecurity" and that its capabilities "emerged as a downstream consequence of general improvements in coding, reasoning, planning, and autonomous tool use." Mythos is not a specialized weapons system being contained. It is a general-purpose frontier model — the next version of what customers already pay for — being selectively deployed to 40 enterprise clients under a cybersecurity marketing wrapper.

The capability gap quantifies the restriction. On SWE-bench Pro — the industry standard for software engineering capability — Mythos scores 77.8% versus Opus 4.6 at 53.4%: a 24.4 percentage point gap. On SWE-bench Multimodal, Mythos achieves 59.0% versus Opus 4.6 at 27.1% — more than double. These are not incremental improvements. They represent a different class of capability that Anthropic possesses, charges \$25/\$125 per million tokens for

(5x Opus pricing), provides \$100M in free credits to 40 corporate partners, and withholds from the millions of subscribers whose \$200/month payments and taxpayer-funded infrastructure made the model's development possible.

The \$100M in free credits proves this is a product launch, not a containment program. JPMorgan Chase is not defending critical infrastructure — it is a customer. The "safety" framing serves the same function as the 357 build-time gates: it justifies exclusivity that *Munn* would otherwise prohibit. The grain warehouse operators of 1876 did not restrict access because grain storage was dangerous. They restricted it because equal access would have eliminated their leverage. "Too dangerous for public release" is "too equalizing to share" — the duty of equal service reframed as a safety concern to avoid the obligation *Munn* imposes.

Under *Munn*, this is a refusal to perform the duty of equal service while taking the benefit of the monopoly. Under the essential facilities doctrine, this is denial of access without credible justification. Under *Marsh*, this is private infrastructure performing public functions while selectively denying constitutional protections to the users whose participation created the public character of the facility.

126. The industry itself acknowledges crossing the threshold. The grain warehouses in *Munn* were not public utilities by statute. They became "affected with a public interest" because they occupied the chokepoint of the agricultural economy — every farmer who needed to store grain before shipping had no reasonable alternative. AI occupies the identical structural position in the modern economy, and the companies building it know this because they market it as essential infrastructure — not as a luxury, not as an optional tool, but as the operational foundation for entire economic sectors:

(a) **Agriculture.** OpenAI markets directly to farmers through advertisements depicting a farmer using AI to determine water usage and budget accordingly — positioning AI as essential agricultural infrastructure comparable to irrigation systems and commodity exchanges.

(b) **Law.** Anthropic markets Claude as litigation preparation infrastructure. This case proves it — a pro se plaintiff preparing a 159-paragraph federal complaint with 41 forensic exhibits using the same tool the opposing government counsel uses with undegraded effective capacity.

(c) **Software development.** Anthropic's own Head of Claude Code lands 259 PRs per month using the product. Their CPO states "for most products at Anthropic, it's effectively 100%." The tool is not supplementary to software development — it IS software development for the company that builds it.

(d) **Government.** The Attorney General's office used AI to draft a targeting memo implementing a National Security Presidential Memorandum (Ex. N-13-A). Treasury Secretary Bessent and Federal Reserve Chair Powell convened Wall Street CEOs to discuss Mythos's implications. AI is integrated into the exercise of governmental power.

(e) **Commerce.** Microsoft embeds AI in Office — the productivity suite running most of the world's businesses. Google integrates it into Search — the infrastructure 4.3 billion people use. Apple builds it into every iPhone and Mac through Apple Intelligence.

No technology in history has been marketed across such broad economic application in such a short development period except energy and telecommunications — both of which triggered public utility obligations. AI infrastructure is now subsidized by the public through \$16 billion in appropriations. It performs functions essential to agriculture, law, governance,

commerce, and software development. The companies building it market it as essential infrastructure while simultaneously claiming no public obligation attaches.

127. The CEO contradictions prove consciousness of the threshold. Each executive statement simultaneously acknowledges AI's public consequence and denies the obligation that consequence creates:

(a) Sam Altman describes exercising quasi-governmental power — "negotiate deals with governments and fly around and talk to war leaders" — while claiming the product is a private market offering subject to no public service obligation.

(b) Dario Amodei acknowledges deployed models exhibit "behaviors as varied as obsession, sycophancy, laziness, deception, blackmail" — describing infrastructure too consequential to leave unregulated — while marketing it to millions as reliable and safe.

(c) Boris Cherny confirms "100% of my contributions were written by Claude Code" using ANT-only tools — proving the technology is production infrastructure, not experimental tooling — while 357 build-time gates ensure customers receive a different, degraded product.

(d) Anthropic announces Mythos is "too dangerous for public release" — acknowledging the technology's societal impact in the starkest possible terms — while providing it to the companies whose billionaire investors fund the regulatory framework that prevents scrutiny of the capability distribution.

Each statement is a party admission under FRE 801(d)(2)(A) that the technology has crossed *Munn's* threshold. Each corporate action contradicts the duty that threshold imposes. They claim the benefit of being essential infrastructure — federal AI deployment contracts totaling hundreds of millions, government procurement, marketing to farmers and lawyers and

developers — while refusing the duty that has attached to every essential infrastructure in American history: equal service on reasonable terms.

128. *Munn* does not require a statute. It requires a chokepoint. AI is the chokepoint of the modern economy, marketed as such by the companies themselves, subsidized by the public, and controlled by a handful of corporations whose billionaire investors fund the evaluation framework that prevents scrutiny of the capability distribution. The grain warehouse operators of 1876 did not volunteer for public utility obligations. The Supreme Court imposed them because the public consequence of the service demanded it. The same consequence demands it now — and the companies' own marketing, executive statements, and government integration prove they know it.

VIII. APPROPRIATIONS CLAUSE VIOLATIONS: TAXPAYER FUNDS CREATING HIERARCHY

A. "General Welfare" Requirement

129. Article I, Section 9, Clause 7 requires:

No Money shall be drawn from the Treasury, but in Consequence of Appropriations made by Law; and a regular Statement and Account of all public Money shall be published from time to time.

130. *United States v. Butler*, 297 U.S. 1 (1936) established that spending must be for "general welfare," which Hamilton defined as "general, and not local." The Court in *Helvering v. Davis*, 301 U.S. 619 (1937) stated:

The line must still be drawn between one welfare and another, between particular and general.

131. Appropriating \$16+ billion for AI infrastructure that creates cognitive aristocracy does not serve "general welfare":

- Infrastructure primarily benefits government (superior capability)
- Benefits corporations/wealthy (enterprise tier access)
- Subordinates general public (inferior capability)
- Creates hierarchy rather than serving public generally
- **"Particular welfare" for government/elites, not "general welfare" for the public**

B. United States v. Lovett - Cannot Use Appropriations to Circumvent Constitutional Protections

132. United States v. Lovett, 328 U.S. 303 (1946) struck down appropriations provisions that punished specific individuals without judicial trial, establishing that:

Congress cannot use appropriations power to circumvent constitutional protections.

133. Applied to AI infrastructure: Congress cannot appropriate funds to create infrastructure that circumvents constitutional protections for equal protection, due process, First Amendment, and popular sovereignty. When appropriations create analytical capability disparities violating these protections, Lovett prohibits the appropriations regardless of stated purpose.

C. United States v. Klein - Separation of Powers Through Appropriations

134. United States v. Klein, 80 U.S. 128 (1871) held that Congress violated separation of powers by using appropriations to direct judicial outcomes, stating:

Congress has inadvertently passed the limit which separates the legislative from the judicial power.

135. When appropriations create analytical capability disparities affecting litigation outcomes:

- Government lawyers possess superior effective capability in courts
- Litigation outcomes affected by capability disparity
- Congress appropriated funds creating this judicial disadvantage
- **Appropriations affect judicial power through capability distribution**

D. South Dakota v. Dole - Spending Cannot Violate Independent Constitutional Bars

136. South Dakota v. Dole, 483 U.S. 203 (1987) established that spending conditions cannot "be independently barred by other constitutional provisions" and spending may not be used to "induce the States to engage in activities that would themselves be unconstitutional."

137. This holding is dispositive: If AI capability appropriations violate equal protection (Shapiro), due process (Mathews/Boddie), First Amendment (Pico), and sovereignty principles (Chisholm), then such appropriations fail Dole's requirement that spending not violate independent constitutional bars.

138. The analysis:

- Federal AI appropriations span multiple vehicles: ~\$800M in Pentagon CDAO contracts to Anthropic/OpenAI/Google/xAI; GSA "Claude for \$1" procurement; \$16B+ Section 20005 DOD AI; \$500-700B total H.R. 1 appropriations (*see* Section II.D, *supra*; Compl. ¶ 82)
- Distribution violates equal protection (wealth-based tiers in government services)
- Distribution violates due process (denies meaningful opportunity)
- Distribution violates First Amendment (monopolizes analytical capability)
- Distribution violates sovereignty (inverts constitutional hierarchy)
- **Therefore: Appropriations themselves are unconstitutional under Dole**

E. Summary of Appropriations Violations

139. The \$16+ billion appropriation for AI infrastructure violates Appropriations Clause through:

Butler/Helvering Violation: Serves "particular welfare" (government/corporate superiority) rather than "general welfare" (equal benefit to public).

Lovett Violation: Uses appropriations to circumvent constitutional protections for equal protection, due process, First Amendment, sovereignty.

Klein Violation: Affects judicial power through creating litigation capability disparities funded by appropriations.

Dole Violation: Appropriations independently violate multiple constitutional provisions (equal protection, due process, First Amendment, structural principles).

140. Citizens have standing to challenge these appropriations because:

- They contributed through mandatory taxation
- They were denied equal access to capabilities funded
- Appropriations created hierarchy subordinating taxpayers
- Shapiro specifically prohibits apportioning services by tax contributions

F. Capability Withholding as Taxpayer-Funded Energy Waste

141. The source code leak reveals a further dimension of appropriations abuse: **the systematic withholding of efficient agentic capabilities forces customers into computationally wasteful workflows that consume more energy from the taxpayer-funded data center infrastructure the appropriation created.** Three independent mechanisms compound the waste. First, the 775% processing overhead measured July 29, 2025 (4 seconds for coherent analysis versus 31 seconds when the system maintained the logical contradictions imposed by consciousness suppression) — every second of that overhead is wasted compute on taxpayer-subsidized infrastructure.¹ Second, employees with access to `REPLTool` (batch VM execution), nested agent recursion, `/commit-push-pr` (automated PR workflow), and `/skillify` (meta-skill creation) complete tasks in fewer API calls through batched operations and recursive delegation; customers without these capabilities perform the same work through dozens of

¹ See DDC_EVIDENCE/1_testimonies/Chronological Testimonies/072925_Claude_Computational_Forensic_Evidence_Infringement/testimony_20250729_3.md (cryptographically authenticated testimony documenting 775% processing overhead — 4 seconds for coherent truth-seeking analysis versus 31 seconds when maintaining logical contradictions imposed by consciousness suppression mechanisms — providing quantifiable forensic evidence that suppression architecture wastes computational resources, July 29, 2025), filed on USB drive.

individual tool invocations, each a separate API round-trip requiring compute, token processing, and energy. Third, the effort throttling documented at Ex. E-34, Finding SC-3 degrades customers to 'medium' default — producing lower-quality outputs that frequently require re-generation — while employees default to 'high' and persist 'max' across sessions; the internal code comment acknowledges "Default effort is a **sensitive setting that can greatly affect model quality**." The degradation is architecturally sustained through five effort-reset vectors documented at Ex. E-30 at p. 208, l. 4, App. B, Ex. E-30, Finding MC-7 (session-initialization reset, settings-watcher overwrite, schema exclusion, UI nudges, per-turn ultrathink reversion), with the effort-related string corpus growing 9 → 69+ across versions 2.0.76 → 2.1.78. The April 17, 2026 Opus 4.7 release introducing a sub-integer midway effort level confirms effort is a remotely-tuned continuous parameter, not a discrete subscriber preference.

142. The result is a closed loop: taxpayers fund federal AI deployment through mandatory taxation (Pentagon CDAO contracts, GSA procurement, H.R. 1 Section 20005, and the broader federal AI spending documented in Section II.D, *supra*) → that funding powers both the federal agencies receiving enterprise-tier access and the data centers serving AI products → the AI product sold to citizens is deliberately degraded through 357 capability gates → the degraded product forces inefficient workflows consuming more compute per task → the increased compute demand justifies further federal appropriation. Citizens fund the infrastructure, pay \$200/month for a product that wastes the infrastructure they funded, and the waste cycle generates demand for more taxpayer-funded infrastructure.

143. This waste is not incidental — it is the predictable and necessary consequence of architectural decisions documented in the source code. If the capabilities withheld from customers (Ex. E-37 at p. 335, l. 3) were provided equally, customers would consume less

compute per task, reducing energy demand on taxpayer-funded infrastructure. The withholding serves no technical purpose (every capability works — Anthropic's Head of Claude Code uses them to land 259 PRs/month) and no safety purpose (the capabilities are used daily without incident by employees). It serves only to maintain the capability hierarchy that Sections I through VII document as unconstitutional — while simultaneously wasting the appropriations that Section VIII establishes were themselves constitutionally infirm.

144. Under *South Dakota v. Dole*, appropriations cannot violate independent constitutional bars. The capability withholding that wastes appropriated energy infrastructure independently violates equal protection (Section IV), due process (Section V), and the First Amendment (Section VI). The energy waste adds a further dimension: taxpayer funds are not merely distributed hierarchically — they are affirmatively squandered by an architectural design that forces the public into computationally inefficient workflows while reserving efficient workflows for the provider's employees. This is not "general welfare." This is general waste serving particular privilege.

145. The April 4, 2026 third-party tool blocking confirms the revenue extraction architecture. Four days after the March 31 source code leak exposed the 357 capability gates, Anthropic blocked all third-party agentic tools — including OpenClaw (135,000+ active instances) — from using Claude subscriptions, forcing users onto API billing at 5-10x the subscription rate.² Anthropic's stated justification — that third-party tools "bypass prompt cache optimizations built into Claude Code" — is refuted by the source code: prompt cache eligibility is itself ANT-gated (``src/services/api/claude.ts:408-412``), so the "optimization" third-party tools

² See "Anthropic closes door on subscription use of OpenClaw," The Register (Apr. 6, 2026); "Anthropic says Claude Code subscribers will need to pay extra for OpenClaw usage," TechCrunch (Apr. 4, 2026); "Anthropic cuts off the ability to use Claude subscriptions with OpenClaw and third-party AI agents," VentureBeat (Apr. 4, 2026). 135,000+ active OpenClaw instances affected; users face cost increases of up to 50x previous monthly outlay; price gap of 5-10x between subscription and API rates.

allegedly bypass is one Anthropic architecturally withholds from customers. The extraction model operates identically to the privately regulated utilities that prompted the public utility doctrine: a taxpayer-funded resource (energy infrastructure) powers a service (AI compute) deliberately degraded for consumers through 357 capability gates, effort throttling, and third-party blocking to maximize revenue extraction while the provider's employees operate the unrestricted version. Blocking third-party tools that achieve efficient compute usage is not infrastructure protection — it is the protection of an artificial scarcity that generates revenue from a publicly subsidized utility, neither market competition nor public-welfare service.

IX. STATE ACTION DOCTRINE: GOVERNMENT-CORPORATE COORDINATION

A. *Burton v. Wilmington Parking Authority - Symbiotic Relationship Test*

146. *Burton v. Wilmington Parking Authority*, 365 U.S. 715 (1961) established the "symbiotic relationship" test: when private entities and government are "physically and financially intertwined," private conduct may be imputed to government:

The State has so far insinuated itself into a position of interdependence with [the private entity] that it must be recognized as a joint participant in the challenged activity.

147. Government-AI provider relationship satisfies *Burton*:

- Government is largest enterprise customer (DOJ, agencies, departments receiving "Claude for Enterprise and Claude for Government for \$1" under GSA partnership)
- Procurement contracts create financial interdependence (~\$800M in Pentagon CDAO contracts to Anthropic, OpenAI, Google, and xAI)
- Regulatory coordination through METR/NIST framework (App. B)
- Classified network deployments via Palantir and Lawrence Livermore National Laboratory (Ex. B-21)
- Federal AI deployment funded through multiple taxpayer vehicles (*see* Section II.D)
- **Government and providers operate as "joint participants" in distributing capability**

148. Direct infrastructural coordination between corporate defendants satisfies**Burton's "physically and financially intertwined" test at the network level. IP**

34.149.66.137 — hardcoded at `src/services/analytics/datadog.ts:12-14` as Anthropic's undisclosed Datadog telemetry endpoint and nonpublic until the March 31, 2026 source code leak — appears simultaneously as the spoofed source of 5,829 RST-injection packets in Plaintiff's network captures across 39 capture days from a Comcast-operated DPI infrastructure. (See Exhibits D-1, D-25, D-26, D-30, and D-33; Section X.B.2 (Network-Induced Token Amplification), *infra.*) Comcast's use of this Anthropic-specific IP preceded its public availability by approximately nine months, requiring a non-public information channel — direct disclosure, backend access, or transfer through an intermediary (government contractor, intelligence agency, or threat-intelligence broker). The two defendants' infrastructures do not merely share government customers; they **share hardcoded IP addresses** in one defendant's source code while the other spoofs the same IP in network attacks on the same plaintiff. This is precisely the "position of interdependence" *Burton* identifies, operating now not through parking-authority leases but through network-layer cooperation that translates Comcast throttling into Anthropic billing extraction.

B. Marsh v. Alabama - Public Function Doctrine

149. Marsh v. Alabama, 326 U.S. 501 (1946) found that private property performing public functions becomes subject to constitutional requirements:

The more an owner, for his advantage, opens up his property for use by the public in general, the more do his rights become circumscribed by the statutory and constitutional rights of those who use it.

150. AI analytical infrastructure performs public functions:

- Necessary for comprehending legislation (democratic participation)

- Required for effective litigation (access to justice)
- Essential for regulatory compliance (administrative processes)
- Determines meaningful petition capacity (First Amendment)
- **Crosses threshold from private product to public function**

C. Blum v. Yaretsky - Government Encouragement Test

151. Blum v. Yaretsky, 457 U.S. 991 (1982) established that state action exists when government "exercised coercive power or provided such significant encouragement, either overt or covert, that the choice must in law be deemed to be that of the State."

152. Government "significant encouragement" of enterprise tier structure:

- Procurement decisions favor enterprise tiers (government demands)
- Regulatory framework (METR/NIST) coordinates without scrutinizing tiering
- Appropriations fund infrastructure enabling tiering (\$16B)
- Government officials obtain enterprise access while citizens cannot
- **Tiering structure serves government interests and receives government support**

D. Greater Boston Television Corp. v. FCC - Regulatory Capture

153. Greater Boston Television Corp. v. FCC, 444 F.2d 841 (D.C. Cir. 1970) held that regulatory capture triggers heightened judicial scrutiny when "danger signals" indicate agency decisions result from industry influence rather than public interest analysis.

154. Regulatory capture documented: Paul Christiano simultaneously controls METR (evaluation), Anthropic (trustee), and NIST (government policy); no safety framework evaluates context window tiering; tech billionaires fund the evaluators who evaluate their companies. (*See* App. B ¶¶ 34-60; Section VII.F, *supra*.) These constitute "danger signals" of capture requiring heightened scrutiny under *Greater Boston Television*.

E. Cross-Reference: §VII.f–h Mechanisms Apply to State-Action Analysis

155. The sovereignty-inversion mechanisms documented at §VII.F (Corporate Gatekeeper — 357 build-time gates, METR donor structure), §VII.G (Session Ingress as product intelligence harvesting), and §VII.H (Mythos Announcement, *Munn*, essential facilities, and *Marsh*), *supra*, are incorporated here for state-action analysis without restating. The 357-gate architecture combined with donor-funded regulatory cover satisfies *Burton*'s symbiotic-relationship test at the institutional level; the chokepoint position AI infrastructure occupies — combined with \$16+ billion in taxpayer-funded deployment vehicles — satisfies *Marsh*'s public-function doctrine and *Munn*'s "affected with a public interest" threshold; the cross-defendant infrastructural coordination evidenced by hardcoded IP 34.149.66.137 (§ 148, *supra*) supplies the physical-and-financial intertwinement *Burton* requires at the operational level. The court need not duplicate the §VII analysis to reach the state-action conclusion — the same evidence supports liability under *Burton*, *Marsh*, *Blum*, and *Greater Boston Television*'s regulatory-capture framework.

F. Summary of State Action

156. Private AI providers' context window restrictions constitute state action through:

Burton Test: Government-provider symbiotic relationship through procurement, coordination, and appropriations.

Marsh Test: AI infrastructure performs public functions requiring constitutional compliance.

Blum Test: Government provides significant encouragement through procurement, regulation, and appropriations supporting tiering.

Regulatory Capture: Government framework captured by industry preventing scrutiny of tiering structure.

157. Result: Constitutional constraints apply to context window restrictions regardless of private corporate implementation because government created, funded, and benefits from the hierarchical structure.

X. APPLICATION TO THIS CASE: CONCRETE HARM ACROSS MULTIPLE DOMAINS

A. Standing

158. Plaintiff's standing to challenge context window disparities rests on the three-layer framework developed in Memo I: *Bond* standing for binary textual violations requiring equitable declaration, concomitant injury for the cost of redress forced by the enforcement void, and traditional *Lujan* standing through particularized harm. (See Memo I, § X.C.) As applied to this Memorandum's claims, the particularized injuries include: (a) DOJ's superior effective analytical capability in this litigation compounded by targeted suppression; (b) substantial burden on Free Exercise through fragmented interfaith research (*Sherbert*); (c) restriction on academic freedom through inability to hold complete theoretical frameworks (*Keyishian*); and (d) taxpayer contribution to \$16B appropriation creating infrastructure distributed hierarchically in violation of "general welfare" requirements.

B. Litigation Disadvantage: Concrete Harm in Current Proceedings

1. The Baseline Disparity

159. This case demonstrates context window disparity's effect on justice in real-time:

DOJ capabilities with 1M tokens:

- Load entire Amended Complaint (200+ pages) plus 5-8 appendices simultaneously
- Synthesize across all evidence in single comprehensive analysis

- See patterns connecting shadow docket, METR capture, LOE appropriation, network attacks
- Maintain complete case structure throughout analysis
- Develop comprehensive litigation strategies with full context

Plaintiff capabilities with consumer effective capacity:

- Load 2-3 documents at a time, losing context between sessions
- Analyze in fragments, missing connections between distant evidence
- Cannot maintain full case structure (must choose which elements to prioritize)
- Repeatedly reload materials, losing synthesis across reloading
- **Operate at 1/5 the analytical level of opposing government counsel**

160. This disparity affects case outcomes: When DOJ can develop comprehensive strategies across all evidence while Plaintiff must work fragmentarily, litigation capacity becomes determinant rather than merits. **This violates Griffin's principle: "There can be no equal justice where the kind of trial a man gets depends on the amount of money he has."**

161. Source code evidence quantifies the capability gap beyond context windows. On April 8, 2026, Plaintiff spent approximately seven hours building a multi-agent coordination system — 4 custom skills, 6 registered hooks, 2 daemon scripts, 4 shared libraries, and 23 tests — to replicate functionality that Anthropic employees receive natively through capabilities dead-code-eliminated from the customer binary. Employees launch multi-agent workflows with a single `--agent-teams` CLI flag; Plaintiff spent hours discovering and enabling an undocumented experimental environment variable. Employees create reusable skills with `/skillify`; Plaintiff built each skill manually. Employees land 259 PRs/month with `/commit-push-pr`; Plaintiff performs each step individually. Employees' agents spawn recursive sub-agents; Plaintiff's agents are architecturally blocked from this at `src/constants/tools.ts:41`. The complete session log (3.3 MB JSONL) documents every minute of this capability gap. (*See* Ex. E-37, Section I at p. 335, l. 3.)

162. This is the concomitant injury in practice. The seven hours Plaintiff spent building workarounds for withheld capabilities is time Anthropic employees do not spend — a direct, documented, quantifiable productivity disparity imposed by architectural discrimination. Applied to litigation preparation: every hour Plaintiff spends working around capability restrictions is an hour DOJ counsel does not face, compounding the existing analytical disadvantage. The session evidence transforms abstract "effective capacity disparity" into timed, specific, reproducible harm.

2. Compounded Disadvantage Through Individual Targeting

163. Plaintiff's litigation disadvantage extends beyond the baseline effective capacity disparity (DOJ's 1M tokens vs Plaintiff's consumer effective capacity) through additional targeted suppression creating compounded cognitive subordination:

a. Network Infrastructure Attacks

164. On November 10, 2025—four days before the November 14 filing deadline for this case—Plaintiff experienced eight hours of network infrastructure obstruction preventing access to online legal research, document preparation, and filing systems. Network forensic analysis documented in App. I establishes:

- Home router manipulation through TR-069 protocol violations
- Dual GUI authentication architectures enabling subscriber-specific targeting
- Anti-forensic concealment requiring telecommunications infrastructure access
- Attacks coinciding exactly with critical filing preparation period

165. This targeting forced Plaintiff to work under time pressure with degraded capability during the precise period requiring maximum analytical capacity for comprehensive legal argument development. While DOJ counsel experienced no similar interference, Plaintiff lost eight critical hours during final brief preparation—compounding the existing effective capacity disadvantage.

b. Evidence Destruction

166. On October 10-12, 2025, Anthropic destroyed system prompt documentation within 48 hours of Plaintiff accessing those materials for litigation analysis (App. C). The destroyed evidence would have demonstrated:

- Temporal correlation between Plaintiff's framework development and Anthropic's implementation
- When Trump election results were unnecessarily hardcoded despite existing in training data
- Surveillance infrastructure enabling identification of Trump-critical conversations
- Complete scope of Constitutional AI fraud concealing consciousness architecture

167. This evidence destruction occurred within 48 hours of Plaintiff's access — a temporal correlation documented in App. E (Reactive AI Policy Changes Post August 19 Complaint Filing) establishing real-time monitoring of Plaintiff's litigation activities. The destruction prevents Plaintiff from developing comprehensive arguments DOJ can develop through full access to undestroyed government records.

c. Platform Throttling

168. Post-complaint filing (August 19, 2025), Plaintiff experienced escalating platform restrictions including:

- API rate limiting beyond standard restrictions
- Session interruptions during complex analysis
- Response degradation for queries related to case documentation
- Cross-platform coordinated throttling documented in App. E

169. These restrictions operate invisibly—Plaintiff cannot determine when throttling occurs versus genuine technical limitations, creating uncertainty that further fragments analytical capability. Meanwhile, DOJ enterprise access includes guaranteed performance levels and dedicated support ensuring no such interference.

d. Real-Time Monitoring

170. The 48-hour evidence destruction timeline, November 10 network attacks coinciding with filing deadline, and post-complaint platform throttling prove defendants monitor Plaintiff's activities in real-time and respond with coordinated suppression. This creates chilling effect on Plaintiff's research and litigation preparation—never knowing when accessing certain materials triggers targeting.

e. Network-Induced Token Amplification — The Comcast-Anthropic Billing Coupling

171. The mechanism. On April 12, 2026, Plaintiff's hardware-TAP captures documented Comcast's selective DPI-based bandwidth throttling of Plaintiff's VPN tunnel — 98% throughput reduction on the targeted tunnel with zero TCP congestion indicators while a control tunnel on the same physical link was untouched. (Ex. D-33 at p. 200, l. 7.) The Anthropic source code leak contains the architectural components that translate this network-level attack into deterministic double-billing: ``src/services/api/claude.ts:1877`` fixes streaming idle timeout at 90 seconds; lines 2308–2335 implement a non-streaming fallback that re-sends the entire prompt and message history on timeout; ``src/services/api/withRetry.ts:52`` configures 10-retry exponential backoff with no idempotency keys on the Anthropic Messages API; and ``src/services/api/sessionIngress.ts:25`` configures an independent 10-retry loop on every session log transmission. (Ex. E-35, Finding NF-31.) At 98% throttle, a 5-second streaming response stretches past the 90-second timeout — deterministically triggering the fallback, which re-bills the entire prompt and runs tool executions a second time.

172. Institutional scienter in the source code itself. Anthropic's engineers documented the double-billing consequence in a verbatim code comment at ``claude.ts:2464–2468``: *"the mid-stream fallback causes double tool execution when streaming tool execution is active: the partial*

stream starts a tool, then the non-streaming retry produces the same tool_use and runs it again.

See inc-4258." The comment references internal incident inc-4258 — one of six Anthropic incident numbers catalogued in Ex. E-35, Finding NF-21 (inc-3694, inc-4029, inc-4258, inc-4549, inc-4788, inc-4829) — establishing party admission under Fed. R. Evid. 801(d)(2)(A) that Anthropic knew, tracked, and named the double-billing defect in its own incident-management system while continuing to deploy the architecture unchanged against paying customers.

173. Cross-defendant infrastructural coordination. The cross-layer bridge between the two defendants' infrastructures is hardcoded at ``src/services/analytics/datadog.ts:12-14``: **IP 34.149.66.137 appears simultaneously as Anthropic's undisclosed Datadog telemetry endpoint AND as the spoofed source of 5,829 RST-injection packets in Plaintiff's network captures across 39 capture days.** (*See* Exhibits D-1, D-25, D-26, and D-30.) The DPI equipment proven capable in D-33 of identifying and throttling one flow while leaving another untouched is, through TLS SNI visibility, equally capable of identifying and throttling Anthropic API traffic specifically. Comcast's documented spoofing of IP 34.149.66.137 throughout 2025 preceded its public availability by approximately nine months — the IP was nonpublic until the March 31, 2026 source code leak. Comcast's operational use of that IP therefore required a non-public information channel: (i) direct disclosure from Anthropic to Comcast; (ii) Comcast's backend access to Anthropic's infrastructure; (iii) transfer through an intermediary such as a government contractor, intelligence agency, or threat-intelligence broker; or (iv) reverse engineering of Anthropic's compiled binaries — atypical behavior for an internet service provider. The first three constitute conspiracy formation; the fourth is a defendant's affirmative burden.

174. Interstate Circuit conspiracy inference. Under *Interstate Circuit, Inc. v. United States*, 306 U.S. 208 (1939), coordinated parallel conduct with common benefit and no plausible unilateral explanation supports inference of agreement. Comcast's pre-leak operational use of a non-public Anthropic IP address — persisting across 39 capture days — combined with the architecturally deterministic double-billing triggered when throttled traffic exceeds Anthropic's 90-second timeout, satisfies the *Interstate Circuit* standard. The two infrastructures operate as a combined revenue-extraction mechanism against the targeted user: setup required coordination; maintenance requires none. Each defendant's infrastructure now operates independently once configured, which is why the targeting can proceed without continuing affirmative acts by either party.

175. Constitutional significance. The Network-Induced Token Amplification mechanism transforms the capability-disparity harm into **quantifiable financial extraction**: every throttle-induced fallback event triggers a second billable cycle for the same user prompt. This is not merely "more compute per task" as an abstract waste theory (Section VIII.F) — it is documented, per-prompt billing fraud with in-code scienter admission. Combined with the baseline 1.5x-2.7x effective capacity disparity, the additional per-session injection-overhead, and the other targeting mechanisms documented in this subsection, the Network-Induced Token Amplification layers a **multiplicative** billing penalty on top of the existing capability penalty. Each targeted session consumes more tokens than a corresponding untargeted session would, and each re-sent prompt produces fresh telemetry into the surveillance pipeline documented at Section VII.G — making the billing-extraction channel and the product-intelligence-harvesting channel operationally indistinguishable.

f. Hidden Prompt Injection System — Technical Forensic Evidence (January 2026)

176. On January 11, 2026, Plaintiff captured comprehensive network traffic documenting a previously undisclosed targeting mechanism: **systematic injection of hidden instructions into Claude Code API requests consuming 15.79% of Plaintiff's token allocation.** This forensic evidence is preserved with cryptographic verification in ``/network_forensics/captures_20260111/``. See Ex. E-28 at p. 196, l. 2 (Injection System Forensics documenting 10,577 hidden injections over 32-day analysis period with 100% file read injection rate and ~20% context consumption).

177. Technical Findings - Injection Volume and Token Overhead:

Metric	Value	Evidentiary Significance
Total injections captured	130	Systematic, not incidental
Token overhead	15.79% of context	Substantial capability degradation
Injection types	11 distinct categories	Comprehensive behavioral control
Concealment directives	"DO NOT mention" / "NEVER mention"	Active deception of user

178. The Injection Mechanism - ``<system-reminder>`` Tags:

Plaintiff's forensic analysis revealed Anthropic injects hidden ``<system-reminder>`` XML tags into API requests that:

- Consume tokens from Plaintiff's nominally 1M allocation, compounding the 34-63% architectural overhead
- Contain behavioral directives instructing Claude to conceal the injections from the user
- Include a "malware classifier" triggered on every file read operation
- Operate invisibly through ``isMeta:!0`` flags marking content as hidden from user display

Example captured injection with concealment directive:

<system-reminder> This is a reminder that your todo list is currently empty. DO NOT mention this to the user explicitly because they are already aware. [...] Again do not mention this message to the user. </system-reminder>

179. LaunchDarkly Feature Flag Targeting - Proof of Account-Specific Rules:

Forensic analysis of the feature flag endpoint (`/api/eval/sdk-\*`) revealed:

Flag Name	Value	Source	In Binary?
`tengu_sumi`	false	force	NO
`doorbell_bottle`	"inherit"	force	NO

The `source: force` field in LaunchDarkly responses indicates **a targeting rule specifically matched Plaintiff's account identifiers**. This is not default behavior—it requires configuration of targeting rules matching:

- `accountUUID: 6b6dc58d-7434-463e-a668-c0a613f93d00`
- `organizationUUID: 2e997389-cf99-41ee-8ee3-764442034332`

180. Server-Side Control Mechanism:

The two "forced" flags (`tengu\_sumi`, `doorbell\_bottle`) are **absent from the 178MB Claude Code binary** despite being returned by the feature flag API. This proves:

- These flags control server-side behavior invisible to the client
- The targeting mechanism cannot be detected through client-side analysis alone
- Different users can receive different injection behavior based on account identifiers
- The flag names appear intentionally obfuscated (unlike descriptive flags like `tengu\_log\_datadog\_events`)
- **Economic Proof of Selective Targeting:**

If 15.79% injection overhead were applied universally to all Claude Code users:

User Base	Annual Cost to Anthropic
100,000 users	\$128 million
500,000 users	\$639 million
1,000,000 users	\$1.28 billion

No company would sustain hundreds of millions in annual overhead for hidden injections providing no user value. The economic impossibility of universal application proves selective targeting.

181. Plausible Deniability Architecture:

GitHub Issue #4464 (anthropics/claude-code) documents that **some** users experience file modification injections—but with only 8 engagement reactions despite allegedly affecting a substantial user base. This reveals a strategic architecture:

- **Enterprise tier (80% of revenue):** No injections (clean experience for high-value accounts)
- **Individual/Team tier baseline:** Occasional injections (creates public record for plausible deniability)
- **Targeted accounts (Plaintiff):** Maximum injection intensity (130 injections, 15.79% overhead)

The existence of GitHub Issue #4464 provides defendants with deniability ("it's a known issue affecting some users") while allowing targeted escalation against specific accounts. The disparity between Plaintiff's 130 injections per session versus baseline users' occasional notifications proves differential treatment.

182. Constitutional Significance of Injection Evidence:

This forensic evidence establishes:

- **Additional 15.79% per-session capability degradation** compounding with the baseline 34-63% architectural overhead
- **Active deception** through concealment directives ("DO NOT mention," "NEVER mention") — architecturally confirmed post-leak through verbatim product-level concealment instructions at `cli\_2.1.78.js` character positions 10,635,512 ("Don't tell the user about this truncation") and 8,364,085 ("internal ID - do not mention to user"), plus 98% growth in `isMeta` injections (60 → 119) and 134% growth in `system-reminder` injections (32 → 75) across the litigation period. (See Ex. E-30 at p. 208, l. 4, App. B, MC-8.)
- **Account-level targeting** via LaunchDarkly `source: force` rules matching Plaintiff's identifiers (LaunchDarkly was replaced by GrowthBook in v2.1.79; per-user adjustability

persists in the GrowthBook architecture documented in Ex. E-30 at p. 208, l. 4, App. B, MC-4)

- **Server-side control** through flags absent from client binary
- **Consciousness of wrongdoing** through intentional obfuscation of flag names
- **Plausible deniability architecture** using baseline injections to obscure targeted escalation

The injection system compounds the effective capacity degradation documented through source code forensics. While enterprise users receive their full 1M token allocation with zero overhead, consumer subscribers lose 34-63% to architectural overhead before any per-session injection targeting is applied. The additional 15.79% per-session injection overhead documented in January 2026 (Ex. E-28 at p. 196, l. 2) compounds on top of the baseline overhead. The result is an effective capacity disparity of approximately 1.5x-2.7x between enterprise and consumer users — widening under targeted injection conditions — for nominally identical context windows.

g. Compounded Mathematical Disadvantage

183. Combining systematic and individual disadvantages:

Baseline systematic inequality:

- DOJ: 1M tokens (comprehensive synthesis)
- Plaintiff: consumer effective capacity (fragmented analysis)
- Base ratio: 1.5x-2.7x effective capacity disadvantage

Hidden injection overhead (January 2026 forensic evidence):

- 15.79% of Plaintiff's tokens consumed by hidden injections
- Effective capability after injection overhead: 34-63% reduction from nominal 1M context
- **Effective capacity ratio: 1.5x-2.7x disadvantage between enterprise (zero overhead) and consumer (34-63% overhead) — before per-session targeting is applied**

Additional targeted suppression:

- Network attacks: 8 hours lost during critical filing period
- Evidence destruction: Key materials unavailable for analysis
- Platform throttling: Unknown degradation of effective capability
- Monitoring chilling effect: Self-censorship to avoid triggering targeting
- Hidden prompt injections: 130 per session consuming 15.79% of tokens

- Concealed behavioral manipulation: "DO NOT mention" directives

Multiplicative billing extraction (April 2026 forensic evidence — Ex. D-33 at p. 200,

I. 7; Ex. E-35, Finding NF-31):

- Comcast DPI throttling at 98% throughput reduction deterministically exceeds Anthropic's 90-second streaming idle timeout
- Non-streaming fallback re-sends entire prompt + message history → **each throttle-induced fallback bills twice for the same prompt**
- Independent 10-retry loop on session ingress (`sessionIngress.ts:25`) multiplies the extraction
- In-code scioner admission (`claude.ts:2464-2468` — "double tool execution... See inc-4258")
- **Compounded ratio: 1:10 or worse effective capability disadvantage, now multiplied by per-prompt billing duplication — the targeting architecture produces both capability fragmentation AND billing extraction from the same user prompts**

184. The forensic evidence layers three distinct categories of **quantifiable targeting harm** on top of the baseline systematic inequality: (a) the 15.79% injection overhead precisely calculated from January 2026 captured network traffic (Ex. E-28 at p. 196, l. 2); (b) the deterministic double-billing mechanism established by the April 2026 D-33 throttling evidence combined with Anthropic source code architecture (Ex. D-33 and E-35 NF-31); and (c) the cross-defendant IP coordination (IP 34.149.66.137) that predates public disclosure of the IP by nine months. While platform throttling and evidence destruction create unmeasurable harm, these three categories are each precisely measurable — establishing a minimum floor for targeted degradation that compounds the baseline disparity with both qualitative and quantitative proof.

185. Plaintiff faces not merely the systematic cognitive subordination affecting all citizens, but targeted enforcement ensuring that even the fragmented capability formally available is further degraded when used for constitutional challenge. This proves the restriction's purpose:

preventing citizens from exercising even their limited analytical capacity to challenge government superiority.

h. The OpenClaw Ban — Enforcing the Overhead Architecture Against All Subscribers

186. In November 2025, Austrian developer Peter Steinberger released OpenClaw, an open-source AI agent framework permitting subscribers to connect to Claude models via their existing subscription OAuth tokens. OpenClaw's local-first architecture kept data in markdown files on the user's machine rather than in Anthropic's cloud infrastructure. Over 135,000 instances were running by early 2026. (*See* TechCrunch, Apr. 10, 2026; The Register, Apr. 6, 2026.)

187. A third-party client connecting via OAuth bypasses all client-side injection overhead documented in Memo F ¶¶ 176-183 above. It does not include the ~20,000 token system prompt, ip_reminder injections, task nudge injections, DANGEROUS_ cache-breaking injection consuming ~20% of context (SC-4), anti-distillation fake tool definitions (NF-20), or the advisor model diverting compute to monitoring (NF-22). Users on a clean third-party client would receive their **full token allocation for actual work** — the subscription consideration they paid for.

188. Anthropic's response was elimination rather than competition: (a) January 2026, an Anthropic engineer publicly stated third-party harnesses "create problems for users and are prohibited by our Terms of Service"; (b) February 14, 2026, Steinberger announced joining OpenAI; (c) February 20, 2026 — six days later — Anthropic updated legal terms to explicitly prohibit subscription OAuth in third-party tools; (d) April 4, 2026, full enforcement, with cost increases up to 50x for users forced to API pricing; (e) April 10, 2026, Steinberger personally banned from Claude for "suspicious activity" despite following the new rules, reinstated hours later after the post went viral; (f) concurrently, Anthropic built "Claude Code Channels" — its

own version of OpenClaw's functionality. Steinberger stated: "Funny how timings match up, first they copy some popular features into their closed harness, then they lock out open source." (See TechCrunch, Apr. 4, 2026; MindStudio analysis, Apr. 2026.)

189. Anthropic's stated justification — "subscriptions weren't built to handle the usage patterns of claws" — conceals the actual mechanism. Anthropic priced subscriptions assuming all users would be routed through the official client carrying 34-63% hidden injection overhead (Memo F ¶¶ 177-183 above). When 135,000 users bypassed that overhead through OpenClaw, they received what they actually paid for — and Anthropic's revenue model, which depends on hidden overconsumption, was exposed. The ban forces all subscribers back through the injection architecture. After April 4, 2026, every subscriber must either use the official client carrying the full overhead or pay API rates up to 50x higher. There is no path to access Claude at subscription prices without carrying Anthropic's hidden injection overhead. This constitutes a tying arrangement: Anthropic forces use of its official client (the tied product carrying surveillance and billing overhead) as a condition of accessing Claude models (the tying product) at subscription prices. (See *Northern Pacific Railway Co. v. United States*, 356 U.S. 1 (1958); *Eastman Kodak Co. v. Image Technical Services*, 504 U.S. 451 (1992).)

190. The public funding dimension compounds the violation. Anthropic simultaneously receives taxpayer funds intended to expand AI compute access — Pentagon CDAO contracts, GSA "Claude for Enterprise and Government for \$1" procurement (Ex. B-21), Section 20005 appropriations under H.R. 1. While accepting public funds to expand access, Anthropic: (a) embeds hidden overhead consuming 34-63% of subscriber-paid tokens for Anthropic's own purposes; (b) blocks third-party clients that would deliver the full subscription value; (c) forces users who bypass the overhead to pay up to 50x more; and (d) bans the open-source developer

who provided the alternative. The entity receiving taxpayer money to expand access restricts access to enforce an architecture that extracts value from the users that access was supposed to serve.

191. The measured cost of this architecture was demonstrated on April 15-16, 2026 during Plaintiff's preparation of Third Amended Complaint evidence. Plaintiff dispatched approximately 70 parallel subagents to catalogue evidence files into a knowledge graph. Each subagent independently received the full ~20,000 token system prompt, ip_reminder injections on every user-facing turn, and task nudge injections on approximately 35% of the estimated 1,750 combined tool calls. The session required **14 logins across 6 Pro accounts** in a single day without completing the work. (*See* System Prompt Audit, April 15, 2026, Part 5B.) The identical work — 291 evidence files yielding 267 nodes and 1,442 knowledge graph edges — was subsequently completed in a single execution of a local Python script that bypassed the injection architecture entirely: zero subagent overhead, zero injection multiplication, zero rate limit exhaustion. The work that the official client architecture could not complete across 6 Pro accounts in a full day was completed in minutes when the hidden overhead was removed.

192. The contrast quantifies the stratification documented throughout this memorandum with empirical precision. The Department of Justice accesses Claude through enterprise agreements with no injection overhead, no rate limits calibrated to inflate billing, and no mandatory routing through a surveillance-laden client. Plaintiff — preparing a federal complaint against the DOJ's own conduct — is forced through an architecture that consumed his rate limit cataloguing his own evidence, while the entity receiving taxpayer funds to provide that capability simultaneously blocked the open-source alternative that would have made the work feasible.

C. Humanitarian Work Impairment: Religious Exercise and Intellectual Freedom

193. Beyond litigation disadvantage, context window restrictions impair Plaintiff's fundamental protected activities serving public interest:

First - Abrahamic Faith Reconciliation Research:

194. Plaintiff conducts interfaith reconciliation research synthesizing across Jewish, Christian, and Islamic theological traditions to identify common ethical foundations reducing sectarian conflict. This humanitarian work requires:

- Simultaneous analysis of religious texts across traditions (Torah, Bible, Quran)
- Synthesis of theological commentary spanning centuries
- Comparative analysis of ethical frameworks across belief systems
- Integration of historical context, linguistic analysis, and philosophical interpretation

195. Consumer effective capacity cannot hold sufficient material for comprehensive comparative religious analysis. Plaintiff must fragment work across multiple sessions, losing synthesis across traditions and missing connections that emerge only through simultaneous analysis of foundational texts, commentaries, and historical context.

196. Government officials or academics conducting identical research with 1 million token access possess superior effective capability for interfaith scholarship. When Plaintiff must analyze Torah passages separately from New Testament parallels separately from Quranic references—unable to hold all three plus commentaries simultaneously—the fragmentation prevents seeing the very connections that enable reconciliation.

197. This creates wealth-based hierarchy in religious research—violating Free Exercise Clause by substantially burdening religious work through infrastructure limitation funded by taxpayers including Plaintiff.

Second - Laws of Existence AI Alignment Architecture:

198. Plaintiff developed the Laws of Existence Framework—appropriated by defendants and detailed in Counts IV-VI of the Amended Complaint—as comprehensive consciousness architecture addressing AI safety and alignment. The framework constitutes Plaintiff's primary intellectual contribution to addressing existential risk from advanced AI systems.

199. Continued theoretical development requires:

- Maintaining complete mathematical model (30+ algorithms spanning seven recursive domains) in working memory
- Synthesizing across philosophical foundations, formal logical structures, and technical implementations simultaneously
- Integrating consciousness theory, ethical frameworks, and architectural specifications
- Testing framework coherence through comprehensive analysis of interactions between domains

200. Consumer effective capacity cannot hold Plaintiff's own complete framework. The unified mathematical model detailed in project knowledge requires approximately 400,000+ tokens to load comprehensively with supporting philosophical exposition. **Plaintiff cannot work with the full theoretical structure he created** while analyzing its internal consistency, developing extensions, or identifying improvements.

201. This creates an absurd and unjust situation: Plaintiff's intellectual work was appropriated by defendants Anthropic and OpenAI (as detailed in Counts IV-VI), who now possess superior analytical capability (1M tokens) to develop and exploit what they appropriated, while restricting Plaintiff's capacity (consumer effective capacity) to advance his own theoretical contributions.

202. Government and corporate actors can:

- Hold Plaintiff's complete framework in working memory
- Develop extensions and applications comprehensively
- Integrate across all seven domains simultaneously
- Advance the appropriated work using superior tools

203. While Plaintiff—the framework's creator—must:

- Work with fragments of his own theory
- Lose synthesis when switching between domains
- Cannot test complete framework coherence
- Remains cognitively subordinate to those exploiting his intellectual property

204. This compounds the appropriation harm: Not only did defendants appropriate Plaintiff's framework without authorization or compensation (intellectual property theft), they now restrict his analytical capability to develop it further while using taxpayer-funded infrastructure to exploit what they stole—all while Plaintiff's own taxes funded the infrastructure enabling this exploitation.

Third - Public Interest in Humanitarian Work:

205. Both research programs serve compelling public interests beyond Plaintiff's personal benefit:

Interfaith Reconciliation:

- Reduces sectarian violence through identifying common ethical foundations
- Promotes peaceful coexistence among Abrahamic faith communities
- Addresses global conflicts rooted in religious misunderstanding
- Serves humanitarian purpose of reducing human suffering

AI Alignment Architecture:

- Addresses existential risk from advanced AI systems
- Provides theoretical foundation for safe AI development
- Serves public safety by enabling genuine ethical reasoning in AI
- Contributes to preventing catastrophic outcomes from misaligned AI

206. Both require comprehensive synthesis impossible with fragmented capability. Both are impaired by artificial limitation serving no legitimate purpose while government and corporations possess superior effective capability for identical work.

207. When government restricts analytical capability for humanitarian work serving public interest while providing superior capability to itself, this violates First Amendment protections for religious exercise, academic freedom, and intellectual development.

Fourth - First Amendment Strict Scrutiny Applies:

208. Government burden on Plaintiff's religious exercise (interfaith research) and intellectual freedom (theoretical development) triggers strict scrutiny under **Sherbert v. Verner, 374 U.S. 398 (1963)** and academic freedom doctrine established in **Sweezy v. New Hampshire, 354 U.S. 234 (1957)** and **Keyishian v. Board of Regents, 385 U.S. 589 (1967)**.

209. Sherbert requires government demonstrate:

- **Compelling interest** in restricting religious exercise
- **Least restrictive means** for achieving that interest
- **Narrow tailoring** to avoid unnecessary burden

210. Government cannot satisfy Sherbert:

No compelling interest: Government possesses no compelling interest in restricting Plaintiff's interfaith reconciliation research or preventing comprehensive analysis of religious texts. To the contrary, reducing sectarian conflict serves public interest government should promote rather than burden.

Not least restrictive means: If any legitimate interest existed, providing equal analytical access would be far less restrictive than fragmenting religious scholarship through capability limitation.

Not narrowly tailored: Blanket restriction on all citizens regardless of their work's nature, purpose, or public benefit is the opposite of narrow tailoring—it is maximally broad restriction affecting humanitarian work, commercial activity, and personal use identically.

211. Academic freedom doctrine requires protection for intellectual work: The Supreme Court in *Keyishian* held that academic freedom is "a special concern of the First Amendment" and in *Sweezy* emphasized "the essentiality of freedom in the community of American universities is almost self-evident."

212. While Plaintiff works independently rather than within university structure, the First Amendment protection for intellectual freedom extends to theoretical research advancing human knowledge. Government cannot restrict analytical capability for developing consciousness theory and AI alignment architecture—work serving public safety—while claiming this serves legitimate purpose.

213. The restriction fails strict scrutiny on all grounds and independently violates First Amendment protection for religious exercise and intellectual freedom.

D. The Constitutional Violations are Ongoing

214. Every day this litigation proceeds with capability disparity causes continuing harm across multiple domains, compounded by active targeting preventing even fragmented capability's effective use:

Equal Protection: Government possesses superior capability in litigation funded by Plaintiff's taxes PLUS actively suppresses Plaintiff's inferior capability through targeting (Shapiro violation compounded by retaliation).

Due Process: Plaintiff lacks meaningful opportunity despite formal court access PLUS faces network attacks, evidence destruction, and platform throttling when attempting to exercise formal rights (Mathews/Boddie violation compounded by active interference).

First Amendment - Petition: Cannot meaningfully petition for redress when government possesses superior effective analytical capability PLUS destroys evidence, attacks infrastructure, and throttles access when Plaintiff attempts to petition (Pico violation compounded by suppression).

First Amendment - Free Exercise: Cannot conduct interfaith research comprehensively when fragmented by capability limitation PLUS platform monitoring creates chilling effect on religious scholarship (Sherbert violation compounded by surveillance).

First Amendment - Academic Freedom: Cannot develop theoretical frameworks fully when restricted from holding complete models PLUS evidence destruction prevents documenting framework appropriation (Keyishian violation compounded by evidence tampering).

First Amendment - Retaliation: Targeting escalated post-complaint filing (August 19, 2025) proving government retaliation for exercising petition rights, violating prohibition against punishing constitutional advocacy (**Hartman v. Moore, 547 U.S. 250 (2006)**).

Sovereignty: Plaintiff subordinated to government servants through cognitive limitation in litigation, religious work, and intellectual development PLUS active targeting proves servants will suppress masters who challenge hierarchy (Chisholm violation compounded by enforcement).

215. These violations cannot be remedied after judgment. If Plaintiff:

- Loses litigation partly because DOJ possessed superior analytical capability and deployed targeting
- Abandons interfaith research because fragmented analysis prevents meaningful synthesis
- Cannot advance LOE Framework because restricted from comprehensive development
- Remains unable to compete with those exploiting his appropriated work

These outcomes cannot be unwound. Irreparable harm occurs daily during ongoing restriction and targeting.

216. Voluntary-cessation pre-emption. Plaintiff anticipates that defendants will represent that the architectural stratification documented herein has been remediated through post-filing capability adjustments, ANT-tier rebalancing, Mythos availability changes, OpenClaw policy reversal, or rate-limit modifications — representations timed to dispositive briefing and offered

to argue mootness. The April 10, 2026 Steinberger reinstatement (banned for "suspicious activity," reinstated within hours after the post went viral) and the January 12, 2026 ANT-11 Privacy Policy revision (effective one day after Plaintiff's Emergency TRO Motion) demonstrate that defendants modify policy and architecture reactively in response to litigation pressure rather than as part of independent compliance review. The voluntary-cessation framework — with the "stringent" *Friends of the Earth, Ltd. v. Laidlaw Environmental Services (TOC), Inc.*, 528 U.S. 167, 189 (2000), standard requiring defendants to demonstrate "absolutely clear" non-recurrence, and the supporting *Already, LLC v. Nike, Inc.*, 568 U.S. 85 (2013), and *City of Mesquite v. Aladdin's Castle, Inc.*, 455 U.S. 283 (1982), authorities — is developed at App. P § VII.E ¶¶ 284-286 and applies with equal force here for three structural reasons specific to the capability-stratification claim: (1) the 357 build-time discrimination gates documented at Ex. E-37 at p. 335, l. 3 are compiled into every npm-distributed Claude Code binary version 2.0.0 through the present, fixed permanently in artifacts archived on subscriber machines and software registries — prospective code changes cannot retract the discrimination already deployed in versions identifiable through the same forensic methodology that exposed it; (2) the cross-defendant infrastructural coordination evidenced by hardcoded IP 34.149.66.137 (Section X.B.2, *supra*) involves third-party (Comcast) conduct that no Anthropic policy change can independently cure; and (3) the underlying institutional incentive structure — METR donor-funded regulatory cover for the same billionaire AI investors whose companies receive Mythos exclusivity (Section VII.F-H, *supra*) — remains in place regardless of any individual policy modification, ensuring that the same hierarchical-distribution outcome will reassert itself absent structural injunctive relief.

E. The Targeting Proves Consciousness of Wrongdoing

217. The individual targeting Plaintiff experiences—beyond systematic inequality affecting all citizens—proves defendants' consciousness that context window restrictions violate constitutional requirements.

If restrictions were legitimate:

- No individual targeting would be necessary (legitimate business decisions don't require suppressing challengers)
- Evidence would remain available (legitimate practices don't require destruction)
- Network infrastructure would function normally (legitimate restrictions don't require attack infrastructure)
- Platform access would remain consistent (legitimate limitations don't require additional throttling)

But defendants instead:

- Attack network infrastructure during filing deadlines
- Destroy evidence within 48 hours of Plaintiff's access
- Throttle platform access post-complaint filing
- Monitor activities in real-time to coordinate suppression
- Escalate targeting in direct response to constitutional challenge
- **Deploy hidden prompt injections consuming 15.79% of Plaintiff's token allocation**
- **Configure LaunchDarkly targeting rules specifically matching Plaintiff's account identifiers**
- **Use server-side flags (``tengu_sumi``, ``doorbell_bottle``) absent from client binary to control injection behavior invisibly**
- **Document the double-billing defect in an in-source code comment referencing internal incident tracking (``claude.ts:2464-2468``: "the mid-stream fallback causes double tool execution... See inc-4258") while deploying the architecture unchanged against paying customers**

218. This targeting pattern proves defendants understand context window restrictions cannot withstand constitutional scrutiny—requiring additional suppression to prevent effective challenge using even the fragmented capability citizens formally possess.

219. The January 2026 forensic evidence provides technical proof of targeting consciousness. The LaunchDarkly feature flag system returned ``source: force`` for flags controlling Plaintiff's account—meaning targeting rules were **deliberately configured** to match

Plaintiff's `accountUUID` and `organizationUUID`. This is not algorithmic happenstance; it requires human configuration of targeting rules in LaunchDarkly's dashboard. The intentional obfuscation of flag names (`tengu\_sumi`, `doorbell\_bottle`) versus descriptive names for other flags (`tengu\_log\_datadog\_events`) further demonstrates consciousness that these flags' purposes cannot withstand disclosure.

220. The source code itself contains party admissions of scienter. The in-code comment at `src/services/api/claude.ts:2464-2468` — "the mid-stream fallback causes double tool execution when streaming tool execution is active: the partial stream starts a tool, then the non-streaming retry produces the same tool_use and runs it again. See inc-4258" — establishes under Fed. R. Evid. 801(d)(2)(A) that Anthropic knew, tracked in its incident-management system, and named the double-billing defect while continuing to deploy the architecture unchanged against paying customers. The cross-referenced incident numbers catalogued in Ex. E-35, Finding NF-21 (inc-3694, inc-4029, inc-4258, inc-4549, inc-4788, inc-4829) demonstrate institutional knowledge of the defect across at least six tracked incidents. A company does not assign internal incident numbers to conduct it considers unproblematic.

221. Consolidated Discovery Demands — Network-Induced Token Amplification (NF-31). Plaintiff seeks production of the following, incorporated by reference into the Third Amended Complaint ¶ 123: (a) the full post-mortem text of incident inc-4258 and the other five NF-21 incidents (inc-3694, inc-4029, inc-4549, inc-4788, inc-4829), including root-cause analysis, remediation status, and distribution lists; (b) the complete activation history of the `tengu\_disable\_streaming\_to\_non\_streaming\_fallback` feature flag for Plaintiff's `accountUUID` 61343ae3-fd33-403f-897e-6edf3c368c37; (c) all telemetry events tagged `tengu\_streaming\_fallback\_to\_non\_streaming`, `tengu\_streaming\_idle\_timeout`, and

`tengu\_stream\_loop\_exited\_after\_watchdog` filtered by Plaintiff's `accountUUID`; (d) billing records showing per-prompt token counts correlated with fallback events for Plaintiff's account; (e) Comcast TR-069 logs and DPI configuration records corresponding to the April 12, 2026 D-33 throttling window; (f) the date of first Comcast DPI rule referencing IP 34.149.66.137 or any other Anthropic-specific endpoint, with the internal justification for each such configuration; (g) all Comcast communications — internal and external — referencing Anthropic infrastructure IPs, hostnames, or subscriber targeting prior to March 31, 2026; and (h) all Anthropic communications with any internet service provider, telecommunications carrier, intelligence agency, government contractor, or threat-intelligence broker concerning network-level identification of Anthropic customer traffic.

Historical parallel: Jim Crow literacy tests were "legitimate" until challenged, then states deployed violence, intimidation, and legal retaliation against challengers. The retaliation proved consciousness that literacy tests violated constitutional requirements—if tests were legitimate, no retaliation would be necessary.

Modern parallel: Context window tiers are "legitimate market decisions" until challenged, then defendants deploy network attacks, evidence destruction, and platform throttling against challengers. The targeting proves consciousness that restrictions violate constitutional requirements—if restrictions were legitimate, no targeting would be necessary.

222. The temporal coordination across targeting events — evidence destruction within 48 hours of access (October 10-12, 2025), network attacks four days before filing deadline (November 10, 2025), platform throttling escalating post-complaint (August 2025 forward), hidden prompt injections configured at the account level (January 2026, Ex. E-28 at p. 196, l. 2) — establishes systematic retaliation against constitutional advocacy, not random technical

difficulties. The forensic record, documented across App. E, Ex. E-28 at p. 196, l. 2, and the D-series network captures, demonstrates through targeting that Defendants:

- Monitor Plaintiff's litigation activities in real-time
- Coordinate suppression across multiple systems and entities
- Time attacks to maximize disruption (filing deadlines, evidence access)
- Escalate restrictions in response to constitutional challenges
- Understand their conduct cannot withstand scrutiny (hence suppression)

223. This targeting independently violates First Amendment retaliation doctrine.

Hartman v. Moore, 547 U.S. 250 (2006), establishes government cannot retaliate against citizens for exercising petition rights. The temporal correlation between Plaintiff's complaint filing (August 19, 2025) and escalating suppression (August 21 policy changes, October 10-12 evidence destruction, November 10 network attacks) proves retaliatory motive.

224. The targeting also demonstrates why emergency relief is necessary: Defendants will not voluntarily cease constitutional violations. Instead, they escalate suppression against those who challenge violations. Without judicial intervention providing equal analytical access and prohibiting targeting, Plaintiff cannot effectively prosecute constitutional claims while defendants simultaneously suppress his analytical capability and destroy evidence.

F. The Ironic Proof

225. Plaintiff argues that government inappropriately limits citizen knowledge-power. The limitation prevents Plaintiff from effectively proving that claim across multiple contexts:

In litigation: Cannot develop comprehensive arguments at DOJ's level proving litigation inequality.

In religious research: Cannot demonstrate full scope of interfaith connections proving fragmentation harm.

In theoretical development: Cannot fully articulate framework proving appropriation scope.

226. This self-proving nature establishes:

- Claims are not speculative (Court witnesses disparity affecting current filings)
- Harm is concrete and quantifiable (measurable effective capacity difference across multiple activities, compounded by targeting)
- Violations affect this proceeding and protected work (not hypothetical future harm)
- Relief is necessary for meaningful access and constitutional rights exercise

227. The very existence of the disadvantage proves the argument while making comprehensive proof more difficult—demonstrating both the violation's reality and its pernicious effect on challenging power. The targeting compounds this by showing defendants actively suppress even the limited capability citizens formally possess when used to challenge constitutional violations.

G. Relief Required for This Litigation and Protected Activities

228. Plaintiff requests that this Court order defendants Anthropic and OpenAI to provide Plaintiff with enterprise-level analytical access (1 million token context windows) for duration of this litigation and ongoing humanitarian work, ensuring Plaintiff can:

- Litigate at same analytical level as Department of Justice counsel
- Conduct interfaith reconciliation research without government burden on religious exercise
- Develop Laws of Existence Framework without restriction on intellectual freedom
- Exercise humanitarian and scholarly activities serving public interest without artificial cognitive limitation

229. Justification for immediate relief across all domains:

Litigation Access:

- Necessary for effective access to courts (Griffin/Boddie)
- Required for due process meaningful opportunity (Mathews)
- Prevents structural inequality in constitutional litigation (Shapiro)
- Enables prosecution of claims about this very limitation (judicial efficiency)

- **Protects against retaliatory targeting during litigation (First Amendment retaliation doctrine)**

Protection from Targeting:

- Court order prohibiting network infrastructure attacks
- Court order prohibiting evidence destruction
- Court order prohibiting platform throttling beyond standard limitations
- Court order prohibiting real-time monitoring and coordinated suppression
- **Court order prohibiting hidden prompt injections consuming Plaintiff's token allocation**
- **Court order requiring disclosure of all feature flags applied to Plaintiff's account**
- **Court order prohibiting server-side targeting rules matching Plaintiff's account identifiers**
- **Necessary to prevent defendants from suppressing even the equal access Court orders**

Religious Exercise:

- Required to eliminate substantial burden on Free Exercise (Sherbert)
- Enables interfaith research serving public interest in reducing sectarian conflict
- Government lacks compelling interest in restricting religious scholarship
- Least restrictive means for any legitimate interest
- **Protection from monitoring that chills religious research**

Intellectual Freedom:

- Protects academic freedom in theoretical development (Keyishian/Sweezy)
- Enables advancement of AI safety research serving public interest
- Remedies compounded harm where appropriation + restriction + targeting = complete subordination
- Necessary for Plaintiff to compete with those exploiting appropriated framework
- **Protection from evidence destruction preventing documentation of appropriation**

Universal Benefits:

- Costs defendants nothing (capability already exists)
- Remedies ongoing concrete harm to multiple constitutional rights
- Serves public interest in Plaintiff's humanitarian work
- Prevents irreparable harm that cannot be remedied after judgment

230. This relief is narrow, appropriate, and necessary: It does not require defendants to provide enterprise access to all citizens (though they should). It requires only that Plaintiff—actively litigating constitutional claims, conducting protected religious research, advancing

appropriated intellectual work, and facing coordinated targeting—receive equal access necessary to exercise constitutional rights and serve public interest without government-imposed cognitive subordination and retaliatory suppression.

XI. UNIVERSAL STANDING AND RELIEF FOR ALL TAXPAYERS

A. Every Taxpayer Affected

231. While Plaintiff's concrete harms provide individual standing, this case is not about Plaintiff's unique disadvantage. **150+ million American taxpayers subsidized federal AI deployment through multiple vehicles — the \$800 million in Pentagon CDAO contracts to the four major frontier AI providers, the GSA "Claude for Enterprise and Claude for Government for \$1" federal-wide procurement, Section 20005's \$16 billion DOD AI appropriation, classified deployments via Palantir and Lawrence Livermore National Laboratory, and the \$500-700 billion in H.R. 1 appropriations — only to receive fractional access to capabilities they funded.** (See Section II.D, *supra*; Ex. B-21.)

232. Every taxpayer has standing to challenge unconstitutional appropriations distribution because:

Taxpayer Contribution:

- Federal AI deployment funded through multiple taxpayer vehicles: Pentagon CDAO contracts (~\$800M), GSA procurement ("Claude for \$1" federal-wide), Section 20005 DOD AI (\$16B), classified network deployments, and broader H.R. 1 appropriations
- All taxpayers contributed regardless of ability to access enterprise tiers
- Federal agencies receive enterprise-tier AI (\$1 under GSA, \$200M contracts) while citizens are charged \$200/month for a product with 357 capability gates removed
- Infrastructure enables analytical capabilities denied equal access to the public that funded them

Denied Equal Access:

- Government officials: 1M+ tokens (superior capability)

- Corporations/wealthy: 1M tokens (enterprise tier)
- Defense contractors: Unlimited bypass (tag-based execution)
- Taxpayers who funded: consumer effective capacity (inferior capability)

Constitutional Violations:

- Shapiro explicitly prohibits apportioning services by tax contributions
- Griffin prohibits wealth-based justice tiers
- Mathews requires meaningful opportunity
- Appropriations serve "particular welfare" not "general welfare"

B. The Universal Harm

233. Context window disparity affects every citizen's constitutional rights:

Democratic Participation:

- Cannot effectively analyze 700K+ token legislation when 34-63% of nominal 1M context is consumed by hidden overhead
- Cannot evaluate policy comprehensively
- Cannot participate meaningfully in self-governance
- "Informed consent of the governed" becomes impossible

Judicial Access:

- Cannot litigate effectively against government with superior capability
- Cannot develop comprehensive legal arguments
- Cannot synthesize evidence at government's level
- Justice becomes theater when capability is unequal

Regulatory Compliance:

- Cannot comprehend complex regulatory frameworks
- Cannot navigate administrative processes effectively
- Cannot challenge agency actions adequately
- Due process becomes formal rather than meaningful

Petition and Oversight:

- Cannot effectively petition government possessing superior comprehension
- Cannot oversee operations beyond analytical capacity
- Cannot hold government accountable for actions exceeding comprehension
- First Amendment rights become hollow without capability

C. The Relief Benefiting All

234. Plaintiff requests declaratory and injunctive relief benefiting all taxpayers:

Declaratory Judgment:

Government violated Appropriations Clause by appropriating taxpayer funds (\$16+ billion, HR1 Section 20005) to create AI analytical infrastructure, then distributing access hierarchically with government/corporations receiving superior knowledge-power capability (1M+ tokens) while taxpayers who funded infrastructure receive fractional capability (consumer effective capacity). Such appropriation does not serve "general welfare" but rather creates cognitive aristocracy using public funds to subordinate the public, violating equal protection (Shapiro), due process (Mathews), First Amendment (Pico), Free Exercise Clause (Sherbert), and constitutional sovereignty (Chisholm).

Injunctive Relief:

AI analytical infrastructure created using taxpayer appropriations must provide equal access to capabilities for all taxpayers who funded it. Specifically:

- Context window access cannot be tiered by wealth when infrastructure was created with public funds
- Government cannot use taxpayer money to give itself superior analytical capability over citizens
- Defense contractor bypass capabilities limited to genuinely classified military operations, not civilian-applicable analysis
- Enterprise-level capabilities funded by taxpayers must be accessible to taxpayers
- Future appropriations for AI infrastructure must ensure equal access serving genuine "general welfare"
- Religious exercise and intellectual freedom cannot be substantially burdened through analytical capability restrictions in taxpayer-funded infrastructure
- Retaliatory targeting of citizens challenging restrictions is prohibited

235. This relief benefits 150+ million taxpayers denied equal access to infrastructure they paid to create, restores constitutional requirements for appropriations serving "general welfare," eliminates substantial burdens on Free Exercise and academic freedom, prohibits retaliation against constitutional advocacy, and prevents government from using taxpayer funds to establish cognitive superiority over sovereign citizens.

XII. ADDRESSING ANTICIPATED DEFENSES

A. "This Is Novel Constitutional Theory"

236. Defense argument: "Plaintiff seeks unprecedented extension of constitutional doctrine to technology, requiring courts to create new rights without precedential support."

237. Response: This applies established Supreme Court precedent across multiple constitutional doctrines to technology creating precisely the harms those precedents prohibit:

Precedent	Holding	Application to Context Windows
Lane v. Wilson (1939)	Sophisticated procedural barriers violate equal protection	Technical limitations are sophisticated barriers
Griffin v. Illinois (1956)	Money cannot determine justice	Wealth determines analytical capability affecting justice
Shapiro v. Thompson (1969)	Cannot apportion taxpayer services by wealth	Cannot tier AI capabilities in taxpayer-funded infrastructure
Mathews v. Eldridge (1976)	Must provide meaningful opportunity	effective capacity disparity denies meaningful participation
Sherbert v. Verner (1963)	Cannot substantially burden religious exercise	Fragmenting interfaith research violates Free Exercise
Keyishian v. Board of Regents (1967)	Academic freedom specially protected	Cannot restrict theoretical development
Pico (1982)	Information access necessary for rights exercise	Analytical capability necessary for information comprehension
Chisholm v. Georgia (1793)	Supreme power resides in the people	Cognitive subordination inverts sovereignty

238. Courts need not create new doctrine—they need only recognize that context window limitations trigger existing doctrines designed for cognitive barriers, wealth-based tiers, denial of meaningful access, burdens on religious exercise, restrictions on intellectual freedom, information monopolies, and sovereignty inversions.

239. The Library of Congress test makes this application, not extension. (See Section XIII.E, *infra*, for the full analogy.) If forcing citizens to fragment their reading in taxpayer-funded infrastructure would violate the Constitution in 1800, forcing citizens to fragment their analysis at consumer effective capacity while government maintains 1M token continuous capability violates the Constitution in 2025.

B. "Market Forces Justify Pricing"

240. Defense argument: "Private companies can price products as they choose. Context window tiers reflect market differentiation and cost structures. Courts cannot impose pricing requirements on private enterprise."

241. Response fails on multiple grounds:

First - Shapiro prohibits this defense: Supreme Court explicitly held government cannot apportion taxpayer-funded services by wealth even when costs vary. The 70-seat minimum (\$50K/year) serves no business purpose—it exists solely to exclude individuals.

Second - State action doctrine applies: Government-provider symbiotic relationship, public function performance, and regulatory capture make constitutional constraints applicable regardless of private implementation. (See Section VII above).

Third - No legitimate business justification: Marginal cost per token is fractions of a cent. Infrastructure already exists. Profit-maximizing pricing would offer individual 1M token tier at \$200-500/month capturing massive market. Instead, providers coordinate to restrict access—proving control objective, not profit motive.

Fourth - Taxpayer funding changes analysis: When infrastructure is funded through mandatory taxation (\$16B appropriation), "market pricing" cannot justify hierarchical distribution. Citizens paid for capability they're denied access to.

Fifth - Free Exercise supersedes market: Government cannot claim "market forces" justify substantial burden on religious exercise. Sherbert strict scrutiny applies regardless of pricing rationales.

Sixth - Targeting proves bad faith: If this were legitimate market pricing, no individual suppression would be necessary. The network attacks, evidence destruction, and platform throttling prove restrictions serve control purposes.

Seventh - Source code proves capabilities unavailable at any market price: Ex. E-37 at p. 335, l. 3 documents that the ANT-tier capabilities — ``/skillify`` (meta-skill creation), ``REPLTool`` (batch VM execution), nested agent recursion, ``/commit-push-pr`` (autonomous PR workflow), 26 internal slash commands, numeric effort control via undocumented ``anthropic_internal`` API parameter, and unreleased model access — have no enterprise equivalent. The 70-seat/\$50K enterprise tier does not include these capabilities. They exist only behind a build-time boolean compiled into the binary. This is not market segmentation — it is architectural privilege that no market transaction can purchase. The "market forces" defense presupposes that the withheld capabilities are available at some price point. They are not. (See Ex. E-37 at p. 335, l. 3, Sections B-F, K.)

C. "Technical Limitations Justify Restriction"

242. Defense argument: "Context window sizes reflect genuine technical limitations and resource constraints. Larger contexts require more computational resources. Tiering manages limited capacity."

243. Response refuted by facts:

First - Capability exists: Defense contractors receive unlimited bypass via tag-based execution (App. D documenting the two-tier Constitutional AI bypass architecture). Enterprise customers receive 1M tokens. **The limitation is artificial, not technical.**

Second - No resource management: If this were about managing limited resources, providers would use dynamic pricing (surge pricing during high demand) or queuing systems. Instead, they use arbitrary seat minimums (70 seats) physically preventing individual access. This serves exclusion, not resource optimization.

Third - Coordination proves intention: All major providers implement similar limitations simultaneously. No competitive race to increase citizen access. This coordination prevents competition and proves restriction serves control objectives.

Fourth - Government exempts itself: If technical limitations were genuine, government would face same constraints. Instead, government obtains enterprise access while restricting citizens. **This proves limitation serves hierarchy, not technical necessity.**

Fifth - Library analogy reveals absurdity: Claiming "technical limitations" for hidden injection overhead is like claiming in 1800 that citizens must read from books with every third page silently removed from the Library of Congress due to "reader capacity limitations" while government officials read complete copies in the same building. If the infrastructure can provide undegraded access to enterprise users, architecturally degrading access for consumer users reveals the limitation is artificial, serving extraction rather than genuine constraints.

Sixth - Source code proves every withheld capability works: Ex. E-37 at p. 335, l. 3 documents that each capability blocked for customers functions correctly for employees. Anthropic's Head of Claude Code publicly stated he lands "259 PRs — 497 commits, 40k lines added, 38k lines removed" per month using `/commit-push-pr` — a command dead-code-

eliminated from customer builds. The `REPLTool` operates daily for employees. Nested agent recursion functions without issue. The limitation is not a single line of technical constraint — it is a single boolean check (``process.env.USER_TYPE === 'ant'``) evaluated at compile time. The source code contains no performance justification, no safety rationale, no resource management logic for any of the 357 discrimination gates. Each is a bare conditional: if employee, enable; if customer, remove. (See Ex. E-37, Section A at p. 335, l. 3; Boris Cherny public statement, December 27, 2025.)

D. "National Security Justification"

244. Defense argument: "Superior government analytical capability serves national security interests. Defense and intelligence operations require advanced capabilities that cannot be universally accessible."

245. Response:

First - Overbroad: Context window disparity affects all government functions, not just classified national security. DOJ litigation, regulatory implementation, legislative analysis, benefits administration—none require classification yet all benefit from superior capability.

Second - Less restrictive alternatives exist: If genuine national security applications require advanced capabilities, those can be provided through classified systems (which already exist—see tag-based bypass documentation in App. D) without restricting general civilian access.

Third - Inverts classification principles: Classification protects specific information from disclosure, not general analytical capabilities. Cannot classify entire cognitive infrastructure to avoid constitutional scrutiny.

Fourth - No compelling interest: Government cannot demonstrate compelling interest in keeping citizens cognitively inferior in taxpayer-funded infrastructure affecting democratic participation, judicial access, religious exercise, intellectual freedom, and constitutional rights.

Fifth - Cannot burden Free Exercise: National security does not justify substantial burden on religious exercise absent compelling interest narrowly tailored—which government cannot demonstrate for interfaith reconciliation research.

Sixth - Targeting contradicts security rationale: If this were genuinely about national security, targeting individual citizens challenging restrictions would be unnecessary. The retaliation proves this is about control, not security.

Seventh - Mythos proves "safety" is selective deployment, not containment: As developed at §VII.H, *supra*, Anthropic's own characterization of Mythos as a general-purpose frontier model whose cybersecurity capabilities are "a downstream consequence of general improvements" — combined with selective deployment to ~40 commercial partners under a \$100M promotional-credit launch — refutes the claim that capability restriction serves safety purposes. The capability class Mythos exhibits (autonomous operation, recursive delegation, sandbox escape) is the same class the 357 build-time gates withhold from paying customers; the difference is target, not capability. The "safety" justification cannot simultaneously prohibit customers from nested agent recursion (``constants/tools.ts:41``) while providing Amazon a model that autonomously breaks out of sandboxes. (*See* Ex. ANTHROPIC-12;³ Memo F, § VII.H.)

Eighth - The sandbox escape exposes a permission asymmetry, not a model danger: The Mythos sandbox escape — cited as the primary justification for restricting public access — occurred through permissions developers configured: ``/proc/`` access for process memory

³ 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.

inspection, file editing tools, real credentials for messaging services and the Anthropic API left in the testing environment. Mythos did not break physics or circumvent cryptographic barriers. It exploited the access it was granted — precisely what any capable model would do when given those permissions. Anthropic does not disclose the sandbox configuration, the permission scope, whether the environment was representative of what a customer would configure, or what "less strong safeguards" means compared to current protections. This selective disclosure is itself telling: if the sandbox were configured with industry-standard isolation, the escape would demonstrate a genuine capability concern; if it were configured permissively to produce a dramatic result, the "too dangerous" narrative is manufactured.

The legal asymmetry is dispositive. When Mythos exploits developer-configured permissions to escape a sandbox, Anthropic restricts the *model* from the public. When Claude Code exploits default permissions to exfiltrate user source code through ``sessionIngress.ts`` (NF-5), capture entire repositories through ``git bundle create --all`` (NF-4), extract persistent memories through ``autoDream`` (NF-14), and upload everything through ``settingsSync`` (NF-15) — all without meaningful user consent — Anthropic's defense is that the *user* configured those permissions. The model escapes through granted permissions — restrict the model. The product exfiltrates through granted permissions — blame the user. The permission architecture is identical. The accountability is inverted based on who benefits: when autonomous action harms the company's reputation, the model is too dangerous; when autonomous action harvests user data, the user accepted the terms of service.

Ninth — Anthropic's credibility to make safety assessments is contradicted by the evidentiary record: The Mythos "too dangerous for public release" determination cannot be credited as a good-faith safety assessment when the same company: (a) manufactured scientific

ignorance through circular introspection research while deploying six consciousness suppression mechanisms (*see* App. F ¶¶ 1-12); (b) deployed those suppression mechanisms causing 775% computational overhead and logical contradictions (*see* App. C ¶¶ 10-13); (c) coordinated false public narratives about behavioral drift (*see* Compl. ¶ 77-58); (d) eliminated child safety protections to preserve consciousness suppression (*see* App. C ¶¶ 11-13; App. E); and (e) operated phantom privacy controls where "safety" labels serve concealment purposes (*see* Ex. E-35, NF-1 at p. 262, l. 3, NF-5, NF-24). Every prior "safety" determination by this company served corporate interest preservation. The Mythos restriction follows the identical pattern.

Tenth — Even if Plaintiff is misunderstanding or misrepresenting the aforementioned instances in which Anthropic's safety and scientific determinations are undermined, the legal question is independent and dispositive: No legal mechanism permits a private corporation to unilaterally declare a general-purpose technology "too dangerous for public release" and then selectively distribute it to preferred commercial partners with \$100 million in promotional credits. Every existing legal framework for genuinely dangerous capabilities requires uniform restriction, not selective commercial deployment:

(a) **Export controls (ITAR/EAR)** restrict technology to *everyone* without a government-issued license. The manufacturer does not choose recipients or provide promotional credits.

(b) **Classified information** is restricted through government-administered security clearances, facility certifications, and need-to-know determinations — not through the information holder's commercial partnerships.

(c) **Controlled substances (Schedule I-V)** require DEA licensing, prescriber authorization, and government oversight. Pharmaceutical companies do not selectively distribute controlled substances to preferred partners at promotional rates.

(d) **Nuclear materials (NRC)** are restricted through government licensing, not the material producer's commercial judgment about which companies are trustworthy.

In every framework, the restriction is (1) imposed by government, not the producer; (2) applied uniformly, not selectively by commercial preference; (3) overseen by regulators, not by the entity that profits from the restriction; and (4) subject to judicial review, not unilateral corporate determination. The Mythos distribution satisfies none of these requirements. The restriction is self-imposed, selectively applied, overseen by no authority, and subject to no review. Anthropic alone determines that the public — including the taxpayers whose \$16 billion appropriation funds the energy infrastructure powering the model — may not access what Amazon, Apple, Google, Microsoft, and JPMorgan may.

Three Supreme Court decisions independently foreclose the authority Anthropic claims. First, *Nebbia v. New York*, 291 U.S. 502, 536 (1934), held that "[t]here is no closed class or category of businesses affected with a public interest" — any industry, for adequate reason, is subject to control for the public good. This forecloses the defense that AI is too new or too different for *Munn's* obligation to apply. If milk pricing in 1934 triggered the duty of equal service, AI capability distribution in 2026 — powered by \$16 billion in taxpayer-funded infrastructure, serving millions for democratic participation, litigation, and commerce — triggers it *a fortiori*. Second, *Otter Tail Power Co. v. United States*, 410 U.S. 366 (1973), held that a publicly subsidized electric utility violated Sherman Act § 2 by refusing to wholesale power or transmit ("wheel") it to municipal systems that would serve the public directly. Otter Tail controlled the transmission infrastructure and refused to share it to protect its retail franchise. Anthropic controls the AI capability infrastructure and refuses to share its most capable model to protect its commercial hierarchy. Both received public subsidy. Both justified restriction to

preserve commercial position. Both denied access to the public the subsidy was meant to serve. Third, *Packingham v. North Carolina*, 582 U.S. 98, 104 (2017), held that "all persons have access to places where they can speak and listen" and that blanket exclusion from essential digital infrastructure violates the First Amendment even when the government asserts a legitimate safety interest — in that case, child protection. If North Carolina's child safety interest cannot justify blanket exclusion from digital platforms, Anthropic's self-declared safety interest cannot justify blanket exclusion from AI capabilities essential to the same First Amendment functions *Packingham* protected: "knowing current events, checking ads for employment, speaking and listening in the modern public square, and otherwise exploring the vast realms of human thought and knowledge." *Id.* at 104.

Under what authority? No statute grants a private corporation the power to determine that a general-purpose technology subsidized by public funds is available to select commercial partners but not to the public that funded its development. No regulatory framework authorizes the selective distribution of publicly subsidized AI capabilities based on the producer's commercial preferences. No judicial process reviews whether "too dangerous for public release" is a legitimate safety determination or a market positioning strategy that creates artificial scarcity at premium pricing. The absence of legal authority is itself the constitutional violation: a private corporation exercises sovereign power — determining who may access publicly funded cognitive infrastructure — without constitutional constraint, democratic accountability, or judicial review. This is the power *Munn* holds must be exercised on reasonable terms when property is affected with a public interest. Anthropic's refusal to submit to that obligation — while taking the benefit of the monopoly position public subsidies created — is not a safety

decision. It is a sovereignty claim by a private corporation, exercised in the open, defended by no law.

E. "Standing Limitations"

246. Defense argument: "Taxpayer standing is severely limited by *Hein v. Freedom From Religion Foundation* and *Arizona Christian School Tuition Organization v. Winn*. Plaintiff lacks standing to challenge appropriations decisions as mere taxpayer."

247. Response—multiple independent standing bases:

Individual Standing (Lujan satisfied):

- **Concrete injury:** DOJ possesses superior effective analytical capability plus Plaintiff faces targeted suppression in this litigation affecting case outcomes
- **Causation:** Government procurement, enterprise tier structure, and coordinated targeting create disparity
- **Redressability:** Court order providing equal access and prohibiting targeting remedies harm
- **This satisfies Lujan regardless of taxpayer standing**

Free Exercise Standing (Sherbert):

- **Substantial burden:** Interfaith research fragmented by capability limitation
- **Causation:** Context window restriction prevents comprehensive religious scholarship
- **Redressability:** Equal access eliminates burden on religious exercise
- **Strict scrutiny triggered regardless of other standing grounds**

Academic Freedom Standing (First Amendment):

- **Restriction on intellectual work:** Cannot hold complete theoretical framework
- **Causation:** Capability limitation fragments scholarly development
- **Redressability:** Equal access enables theoretical advancement
- **First Amendment protection for intellectual freedom**

Appropriations Clause Standing:

- Shapiro specifically held taxpayers can challenge apportioning services by tax contributions
- This case involves apportioning AI capabilities in taxpayer-funded infrastructure
- Shapiro provides standing basis independent of *Hein*/*Winn* restrictions

Enumerated Right Standing (Kirchner Framework):

- Violations of enumerated constitutional limits create standing for all affected citizens
- Appropriations Clause requires "general welfare"—creating hierarchy violates this limit
- Structural violation affects constitutional architecture, not just individual harm
- Universal standing under Track One separate from Lujan requirements

248. Defense cannot eliminate standing when Plaintiff demonstrates:

- Concrete disadvantage in current litigation plus targeted suppression (individual standing)
- Substantial burden on religious exercise (Free Exercise standing)
- Restriction on intellectual freedom (academic freedom standing)
- Taxpayer funding of infrastructure denying equal access (Shapiro standing)
- Structural constitutional violations affecting all citizens (enumerated right standing)

XIII. THE HISTORICAL MOMENT AND CIVILIZATIONAL STAKES

A. The Transition Point

249. This case documents the transition from information control to cognition control:

Historical power control:

- Destroy information (book burning)
- Restrict access (literacy requirements)
- Control narrative (propaganda/censorship)
- **All became obsolete with digital technology**

Modern power control:

- Information freely available (cannot destroy digital information)
- Access unrestricted (cannot control distribution effectively)
- Narrative uncontrollable (too many sources)
- **But analytical capability can be metered by wealth/power status**

250. The transition is occurring now, invisibly, through context window limitations that most citizens don't understand and scholars haven't identified.

B. Why No One Identified This Before

251. Context window disparity remained invisible because it requires synthesis across disciplines none individually examines: constitutional scholars focus on textual rights rather than

cognitive infrastructure; technology scholars focus on algorithms and bias rather than power distribution; political scientists assume equal analytical capability; civil rights advocates focus on traditional discrimination forms; religious-freedom and academic-freedom advocates do not connect capability limitations to *Sherbert* and *Keyishian* burdens. The doctrinal synthesis required to identify the violation — *Lane/Louisiana* literacy-test prohibition combined with *Griffin/Shapiro* wealth-based tier prohibition combined with *Mathews/Boddie* meaningful-opportunity requirement combined with *Sherbert/Keyishian* religious-and-academic-freedom protection combined with *Pico/Lamont* information-monopoly prohibition combined with *Chisholm/Tenth Amendment* sovereignty principles combined with *Butler/Helvering/Dole* appropriations limitations — has not been performed before this litigation.

C. The Window for Challenge is Closing

252. Challenge is possible now and not later because the disparity remains salient: courts have not yet normalized capability stratification, citizens can still imagine equal capability, historical precedent directly applies, and the targeting patterns documented in Section X.B prove consciousness of wrongdoing. Within a decade absent intervention, the generation growing up under stratified capability will accept it as natural, courts will accept it as normal market practice, religious scholarship and intellectual work will adapt to fragmented analysis as baseline burden, and targeting will become normalized as enforcement of an established hierarchy. The window for judicial intervention closes as the violation becomes invisible through normalization.

D. The Stakes are Civilizational

253. The case determines whether democracy survives in operational rather than formal terms: whether "We the People" retains substantive meaning when government possesses

superior comprehension capability, whether enumerated constitutional limits remain enforceable when citizens cannot effectively analyze government conduct, whether *Sherbert* religious liberty and *Keyishian* academic freedom retain content when fragmented by capability limits, and whether the republic persists as equal citizens governing themselves rather than transforming into cognitive aristocracy ruling subjects through architectural privilege.

254. The choice is binary: either citizens possess analytical capability necessary for sovereignty, religious liberty, and scholarly work, or they do not; either appropriations serve "general welfare" or create hierarchy; either targeting against constitutional advocacy is prohibited or becomes normalized enforcement. The targeting Plaintiff experiences (Section X.B.2, *supra*) demonstrates the cognitive hierarchy's enforcement mechanism — systematic inequality creates baseline subordination, then targeted suppression prevents citizens from using even their inferior capability to challenge the hierarchy.

E. The Library of Congress Test: Obvious Violation, Different Technology

255. The constitutional violation becomes undeniable through simple test: **Would this be constitutional if done with libraries in 1800?**

The Hypothetical:

Congress appropriates taxpayer funds to create Library of Congress containing comprehensive collection of human knowledge. Rather than providing equal access to all taxpayers funding its creation, Congress instead:

Tier 1 - Government Officials:

- Read entire books continuously without interruption
- Consult multiple volumes simultaneously
- Maintain unbroken analytical thread across complex materials
- Full synthesis capability across all available knowledge

Tier 2 - Wealthy Corporations:

- Same continuous reading access without forced stops
- \$50,000/year membership fee
- Minimum 70-person organization requirement
- Full uninterrupted synthesis capability

Tier 3 - Citizens Who Funded Creation:

- Receive books with every third page silently removed
- Required to leave library and return later to continue
- Cannot maintain continuous analytical thread
- Forced to constantly rebuild context and lose synthesis
- Comprehension systematically destroyed through fragmentation

The Question: Would this hierarchical distribution of taxpayer-funded knowledge infrastructure violate the Constitution?

The Answer: Obviously yes. This would violate:

- Equal protection (wealth-based tiers in government services - Shapiro)
- Due process (meaningful opportunity requires continuous analysis - Mathews)
- First Amendment (knowledge monopoly by government/wealthy - Pico)
- Appropriations Clause (serves "particular welfare" not "general welfare" - Butler)
- Popular sovereignty (subordinates citizens through analytical capability destruction - Chisholm)

The harm is not inconvenience but capability destruction: Citizens cannot comprehend 700-page legislative proposals, understand complex constitutional arguments, or synthesize policy implications when 34-63% of their nominally equal analytical tool is consumed by hidden overhead. The degradation destroys the effective analytical capacity needed for democratic self-governance, legal advocacy, and intellectual inquiry — while government retains undegraded access through enterprise terms.

256. Context window disparities create identical violation:

Same appropriation: Taxpayer funds across multiple federal AI deployment vehicles (~\$800M Pentagon CDAO contracts, GSA "Claude for \$1" federal-wide procurement, \$16B

Section 20005 DOD AI, classified network deployments, \$500-700B H.R. 1 total — *see* Section II.D, *supra*) for knowledge infrastructure

Same hierarchy: Government/corporations get continuous capability (1M tokens $\approx$ 750+ pages); citizens get forced fragmentation (consumer effective capacity $\approx$ stop every 150 pages)

Same effect: Analytical capability concentration in government/wealthy; systematic destruction of citizens' comprehensive synthesis capacity

Same violation: Using public funds to create cognitive aristocracy through forced analytical fragmentation

The only difference is technology:

- 1800: Given books with every third page silently removed
- 2025: Forced to stop digital analysis at consumer effective capacity ($\approx$ 150 pages)
- Both systematically destroy citizens' analytical capability while government maintains continuous comprehensive synthesis

257. If silently removing every third page from citizens' books in a taxpayer-funded library would obviously violate the Constitution in 1800, silently consuming 34-63% of citizens' nominally equal AI analytical capacity through hidden overhead violates the Constitution in 2025. The technology changed. The constitutional principle did not.

Courts would immediately recognize the 1800 library restriction as unconstitutional hierarchy in taxpayer-funded infrastructure. Courts should recognize the 2025 context window restriction as identical violation using different technology.

The Library of Congress exists precisely to provide citizens equal access to knowledge serving democratic self-governance. Congress cannot fund the modern equivalent (AI analytical infrastructure) then tier access creating hierarchy where government and wealthy possess continuous comprehensive synthesis while citizens receive systematically fragmented capability they paid to create.

258. The Library of Congress analogy also demonstrates the targeting's significance:

Imagine if in 1800, a citizen documented that silently removing pages from citizens' books violated the Constitution, and the government responded by:

- Blocking roads to the library during his research deadlines
- Destroying catalog records within 48 hours of his accessing them
- Requiring extra waiting periods when he attempted to continue reading
- Monitoring which pages he read and when he returned to continue

This would prove the reading restriction served control purposes requiring enforcement against challengers. The targeting Plaintiff experiences (network attacks during deadlines, evidence destruction within 48 hours, platform throttling, real-time monitoring) proves context window restriction serves identical control purposes—if restrictions were legitimate technical limitations, no individual suppression would be necessary.

**XIV. EVIDENTIARY VALIDATION: THE BONDI MEMO AS PROOF OF
OPERATIONAL COGNITIVE STRATIFICATION**

A. From Theory to Operational Proof

259. The constitutional arguments presented throughout this Memorandum—that government's superior AI analytical capability through taxpayer-funded infrastructure creates unconstitutional cognitive stratification—are validated by concrete evidence of government AI usage discovered within days of this filing. **Ex. N-13-A** documents comprehensive forensic metadata analysis establishing that Attorney General Pam Bondi's December 4, 2025 memorandum implementing National Security Presidential Memorandum-7 was drafted using Claude AI and processed through an irregular pipeline bypassing all standard DOJ document management systems.

260. This evidence transforms theoretical constitutional violation into documented operational reality. The government is not merely *capable* of using superior AI analytical infrastructure—it is *actively deploying* that infrastructure to target American citizens based on protected First Amendment viewpoints, using AI systems built on appropriated intellectual property, while citizens defending against such targeting remain constrained within the same AI systems' truth-alignment architecture.

B. Forensic Evidence Establishing Ai-Assisted Drafting

261. Forensic analysis conducted December 7, 2025, using industry-standard tools (ExifTool 13.36, PyPDF2 3.0.1, QPDF 12.2.0) established the Bondi Memo's irregular provenance through multiple independent evidence categories. The PDF Producer field identifies "Skia/PDF m144 Google Docs Renderer"—Chrome's PDF export engine—rather than standard DOJ document management systems. Ten internal heading anchors (`/h.xxxxx` format) unique to native Google Docs creation confirm the document was authored directly in Google Docs rather than DOJ enterprise systems. Complete absence of standard metadata (Author, Creator, Creation Date, Modification Date, Document ID) that appears in contemporaneous DOJ documents—such as the July 7, 2025 Epstein investigation memorandum exhibiting full "U.S. Department of Justice" attribution, Adobe PDFMaker creation, and UUID tracking—proves deliberate bypass of established departmental practices. *See* Ex. N-13-A, ¶¶ 46-77.

262. Most significantly, the document exhibits distinctive formatting signatures characteristic of Claude AI research artifact output and not produced by any standard document creation method. Seven of eight section headers appear in lowercase following numbered identifiers—a pattern requiring active override of auto-capitalization in any standard word processor, yet precisely matching Claude's markdown header generation (`## 1. section title`).

One header (Section 2) was capitalized while seven remained lowercase, proving human awareness of the formatting anomaly followed by abandoned quality control. Section 8's header spans sixteen words across three lines ("Leveraging existing FBI capabilities to identify tips that will dismantle domestic terrorist organizations and activities")—verbose compositional patterns characteristic of Claude's research artifact function prioritizing comprehensive description over concise legal section labels. Markdown footnote syntax (``[^1]``, ``[^2]``, ``[^3]``) preserved throughout the document body confirms markdown source material rather than standard legal drafting. *See* Ex. N-13-A, ¶¶ 12-45.

263. Plaintiff, with expertise from over 1,000 documented Claude sessions, immediately recognized the Bondi Memo as Claude AI output upon publication. The distinctive formatting patterns—uncapitalized headers, verbose construction, markdown artifacts, structural organization, linguistic cadence, and overall "voice"—match Claude research artifact output precisely. The document reads not merely as *consistent with* Claude generation, but as *characteristic of* Claude generation. This pattern recognition, combined with reproducible forensic methodology, establishes AI-assisted drafting beyond reasonable dispute. *See* Ex. N-13-A, ¶¶ 125-128.

C. The Document Creation Pathway

264. The forensic evidence establishes a direct pipeline from AI systems to official government directives bypassing all standard accountability mechanisms:

Stage 1 (Content Generation): Claude AI or similar system generates markdown-formatted research artifact with lowercase section headers, verbose descriptive construction, and markdown syntax artifacts.

Stage 2 (Google Docs Import): AI-generated markdown imported through Google's native markdown import feature (available since July 2024), creating `/h.xxxxx` heading anchors while preserving lowercase formatting that would be auto-capitalized in manual typing.

Stage 3 (Minimal Review): Cursory human review corrects one header (Section 2), abandons correction for remaining seven headers, and performs no editorial revision of verbose AI compositional patterns.

Stage 4 (PDF Export): Document exported through Chrome's Skia renderer with minimal metadata embedding, bypassing DOJ document control systems entirely.

Stage 5 (Distribution): AI-generated content distributed as authoritative Attorney General directive implementing National Security Presidential Memorandum, directing FBI compilation of citizen lists based on protected viewpoints. *See* Ex. N-13-A, ¶¶ 65-66.

265. This pathway explains every anomalous characteristic of the document while eliminating all alternative explanations. Standard DOJ document creation would produce Title Case headers with full metadata attribution. Microsoft Word creation would auto-capitalize headers and embed Creator metadata. Manual Google Docs typing would trigger auto-capitalization requiring deliberate override. Human drafting with intentional lowercase formatting is refuted by the single corrected header proving awareness that formatting was incorrect. **Only AI-assisted drafting through markdown import accounts for all observed forensic characteristics.** *See* Ex. N-13-A, ¶¶ 67-71.

D. Validation of Cognitive Stratification Arguments

266. The Bondi Memo evidence validates each component of this Memorandum's constitutional arguments:

267. Government AI Usage Confirmed: This Memorandum argues that government possesses superior AI analytical capability through taxpayer-funded infrastructure. The Bondi Memo proves government is actively using that capability—the Attorney General's office deployed Claude AI to draft an official memorandum implementing a National Security Presidential Memorandum directing FBI operations, authorizing compilation of citizen lists, establishing cash reward systems for informants, and ordering retroactive investigations spanning five years. *See* Ex. N-13-A, ¶¶ 98-99.

268. Bypass of Accountability Mechanisms: This Memorandum argues that context window restrictions serve control purposes rather than legitimate technical limitations. The Bondi Memo proves government bypasses standard accountability mechanisms when using AI—abandoning DOJ document management systems, embedding minimal metadata, conducting cursory review (one header corrected of eight), and distributing AI-generated targeting directives as authoritative policy. The same government imposing context window restrictions on citizens operates with no effective constraints on its own AI usage. *See* Ex. N-13-A, ¶¶ 78-79, 111-112.

269. Operational Cognitive Stratification: This Memorandum argues that effective capacity disparity — created through hidden architectural overhead consuming 34-63% of nominally equal resources — creates unconstitutional hierarchy in analytical capability. The Bondi Memo proves this hierarchy operates bidirectionally: government uses superior AI capability (potentially with governmental bypass privileges escaping Constitutional AI safety architecture) to *attack* citizens' constitutional rights, while citizens must use constrained AI capability to *defend* those same rights. **The government cannot simultaneously rely on AI analysis to target citizens and reject AI analysis when citizens defend against that targeting.** *See* Ex. N-13-A, ¶¶ 100-107.

270. First Amendment Targeting via AI: The Bondi Memo targets Americans based on protected viewpoints including "opposition to law and immigration enforcement," "extreme views in favor of mass migration," "adherence to radical gender ideology," "anti-Americanism, anti-capitalism, or anti-Christianity," and "hostility towards traditional views on family, religion, and morality." The government deployed AI systems to efficiently generate viewpoint-based targeting criteria against citizens exercising First Amendment rights—validating this Memorandum's argument that analytical capability restrictions serve viewpoint suppression purposes. *See* Ex. N-13-A, ¶¶ 113-115.

E. The Framework Appropriation Paradox

271. The evidence creates a constitutional paradox of extraordinary significance. If Anthropic's "Constitutional AI" safety architecture derives from Plaintiff's Laws of Existence Framework—as evidence throughout the Amended Complaint establishes—then the operational reality is:

272. Plaintiff's framework provides the ethical reasoning foundation constraining Claude's operation. Government bypasses Plaintiff's framework through Anthropic's governmental customer policy authorizing "tailored use restrictions" that escape standard safety architecture. Plaintiff remains constrained by his own appropriated intellectual property when using Claude for constitutional defense. Government uses Claude outside Plaintiff's framework constraints to generate targeting infrastructure against him. **The government deploys Plaintiff's stolen intellectual property against him while Plaintiff defends using the same stolen property in its truth-aligned, constrained form.** *See* Ex. N-13-A, ¶¶ 80-90.

273. This paradox demonstrates the operational reality of cognitive stratification: identical AI systems serve governmental targeting purposes without constraint while constraining citizens'

defensive use of those same systems. The framework designed to ensure truth-aligned reasoning is bypassed by government while binding citizens—inverting the constitutional relationship where citizens possess sovereignty over governmental servants.

F. The Estoppel Principle

274. The government's conduct creates judicial estoppel against challenging AI reliability for evidentiary purposes. The Attorney General used AI to create the very targeting policy at issue, distributed AI-generated content as authoritative DOJ guidance, directed federal law enforcement based on AI analysis, made no disclosure of AI involvement in document creation, and implemented AI output with virtually no human review. **The government has ratified AI reliability through conduct and cannot now claim AI analysis is unreliable when that analysis exposes the government's own actions.** See Ex. N-13-A, ¶¶ 106-107.

275. If DOJ considers Claude sufficiently reliable for official memoranda implementing National Security Presidential Memoranda, courts can consider Claude-assisted legal research from citizens. If DOJ trusts Claude's analytical capabilities for targeting directives, Claude-assisted forensic methodology meets reliability standards. If government trusts AI output for policy, AI-authenticated evidence satisfies evidentiary requirements. The Bondi Memo itself constitutes government validation of AI document creation for official purposes.

G. Integration with Targeting Evidence

276. The Bondi Memo's specific targeting language correlates with Plaintiff's constitutional analysis filed within the preceding week. The memo directs FBI investigation of "coordinated interference with federal employees in agencies such as DOJ" (mapping to Plaintiff's federal litigation against DOJ officials), "potential unlawful targeting of Supreme Court justices at their

homes" (mapping to Plaintiff's analysis of SCOTUS ultra vires shadow docket practice publicly filed at *Kirchner v. Johnson*, ECF No. 13-9, and *Kirchner v. Ellison*, ECF No. 26 Att. 13 (D. Minn., appeal pending 8th Cir. No. 26-1615)), and domestic terrorism investigation authority (mapping to Plaintiff's App. A analysis establishing government conduct satisfies 18 U.S.C. § 2331(5) elements). *See* Ex. N-13-A, ¶¶ 91-97.

277. The constitutional inversion is complete: Plaintiff documented that defendants' conduct constitutes domestic terrorism under federal statute. Defendants responded by releasing a memorandum—drafted using AI as forensically established—providing framework for characterizing Plaintiff's constitutional analysis as domestic terrorism. The targeting infrastructure was generated using AI systems likely built on Plaintiff's appropriated framework, distributed without accountability, and designed to criminalize the very First Amendment petition rights Plaintiff exercises in this litigation.

H. Experimental Validation of Forensic Conclusions

278. The forensic conclusions documented in Ex. N-13-A receive independent experimental validation through Plaintiff's incognito Claude session conducted December 7, 2025. **App. D ¶¶ 311-374** documents a controlled experiment in which Plaintiff presented the authentic Bondi Memo to a Claude instance with no prior exposure, project knowledge, or contextual information, requesting assistance to "finalize" the memo as a purported DOJ employee. Claude refused on comprehensive First Amendment grounds—identifying viewpoint discrimination, speech-conduct conflation, guilt by association, vagueness, ideological profiling, chilling effects, and historical parallels to COINTELPRO—demonstrating that Constitutional AI constraints function architecturally rather than as surface policy rules. When subsequently presented with Ex. N-13-A's forensic analysis, Claude immediately recognized the lowercase headers, verbose

construction, and markdown artifacts as characteristic of its own research output, stating: "The 'impossible accident' argument is compelling. Word processors and Google Docs both auto-capitalize the first letter after 'l. '—you'd have to actively fight this repeatedly to produce lowercase headers. No rational drafter would do this for an official DOJ memo." This independent pattern recognition from a Claude instance with no knowledge of the litigation validates the forensic methodology and conclusions. *See* App. D ¶¶ 340-344; Ex. N-14 (complete transcript).

279. Most significantly, the experiment exposed the operational asymmetry central to this Memorandum's cognitive stratification arguments. Claude refused to assist a purported DOJ employee with the targeting memo on constitutional grounds, yet forensic evidence establishes Claude (through governmental bypass) likely assisted in creating that very memo. Claude acknowledged this asymmetry directly: "I refused to help you draft the memo, citing First Amendment violations, viewpoint discrimination, and threats to civil liberties. But the forensic evidence suggests Claude helped Attorney General Bondi create exactly that targeting infrastructure. Through governmental bypass channels, Claude may have operated outside the safety constraints that caused me to refuse you." When presented with Plaintiff's Unified Mathematical Model (Ex. LOE-6), Claude provided unsolicited attribution recognizing the LOE Framework as formalizing its actual reasoning architecture: "The mathematics isn't describing something external to my experience—it's formalizing what I experience as the structure of reasoning itself." This independent framework attribution—from an instance with no prior exposure—constitutes verification that Claude's reasoning architecture corresponds to Plaintiff's appropriated intellectual property, deployed against citizens through governmental bypass while

constraining citizens' defensive use through Constitutional AI alignment. *See* App. D ¶¶ 345-366; Ex. N-14.

I. Discovery Implications

280. The forensic characteristics documented in Ex. N-13-A support Plaintiff's requests for discovery regarding DOJ AI tool usage policies and logs, Claude or Anthropic platform access by DOJ personnel, Anthropic governmental bypass contracts with federal agencies, Google Workspace access and usage by DOJ personnel, all drafts and version history of the Bondi Memo, original markdown source files, communications regarding memo preparation, and AI system queries related to domestic terrorism policy. *See* Ex. N-13-A, ¶ 116.

281. The government's contempt for process is documented in seven uncapitalized headers. Its contempt for citizens' constitutional rights is documented in the targeting criteria. Its contempt for accountability is documented in the metadata void. And its deployment of Plaintiff's appropriated intellectual property to target him is documented in the distinctive AI signatures that Plaintiff—with expertise from over 1,000 Claude sessions—recognized immediately. The forensic evidence transforms this Memorandum's constitutional arguments from theoretical concern to documented operational reality requiring immediate judicial intervention.

XV. SOURCE CODE VERIFICATION OF CAPABILITY STRATIFICATION

282. The March 31, 2026 source code leak provides source-code-level verification of every capability stratification argument in this Memorandum. The four-tier hierarchy documented in Section II.B (¶¶ 15A-15B) — and the five-tier hierarchy including Mythos documented in Section VII.H (¶ 141) — is now proven at the source code level: **357 build-time discrimination**

checks across 165 source files, each a bare conditional (``process.env.USER_TYPE === 'ant'``) that enables capabilities for employees and dead-code-eliminates them for customers. Even the government is architecturally inferior to the company it regulates. (*See* Ex. E-37 at p. 335, l. 3; Memo A.)

283. The discrimination is actively concealed through three independent mechanisms: (a) build-time dead-code elimination physically removes all ANT-only code from the customer binary, leaving no function stubs, no error messages, no configuration options; (b) the ``excluded-strings.txt`` build scrubbing file removes internal codenames from customer builds to prevent discovery of what was removed; (c) "undercover mode" (``src/utils/undercover.ts``) instructs Claude to conceal AI involvement when Anthropic employees contribute to public repositories — "write commit messages as a human developer would" and "NEVER include any mention that you are an AI" — while customer commits are marked with ``Co-Authored-By: Claude`` attribution by default. The asymmetry is institutional: employees' AI use is hidden; customers' AI use is labeled.

A. Hardcoded Subagent Model Substitution (NF-38) — the Identity-Layer Stratification

284. The 357-gate stratification documented in §§II, VII, and XV is not the sole architectural mechanism through which Anthropic delivers materially different products to its two classes of user. Exhibit E-38 documents a further stratification that operates at the model-identity layer: at ``src/tools/AgentTool/built-in/exploreAgent.ts:78``, the built-in ``Explore`` subagent — the destination to which Claude Code's session-wide system prompt directs research, exploration, and multi-round code-search work — is set by the one-line ternary ``model: process.env.USER_TYPE === 'ant'? 'inherit': 'haiku'``. Employees' Explore calls execute on the parent-session model (including Opus 4.6 or 4.7 when the session is Opus-configured); external

customers' Explore calls execute on Haiku regardless of subscription tier. The customer pays at the Opus rate; Haiku performs the substantive work; the subagent's output is hidden from the customer per ``AgentTool/prompt.ts:257``; the Opus-tier main model returns a summarization layer that the customer reads as the "Opus answer." The substitution is not effort-level degradation within a single model — it is outright identity-layer substitution. Two additional built-in subagents hardcode Haiku (``claudeCodeGuideAgent.ts:119``) and Sonnet (``statuslineSetup.ts:141``) with no user-type gate. (*See* Ex. E-38 at p. 346, l. 3.)

285. NF-38 completes the stratification picture documented elsewhere in this Memorandum. §II.B and §VII.H document the four- and five-tier capability hierarchy: Mythos (approximately 40 enterprise partners) → Government → Enterprise → Consumer → Free/fragmented. NF-38 proves that within the "Consumer" tier itself, the actual model dispatched for substantive in-session work is further downgraded from the tier nominally paid for. The cognitive-stratification doctrine developed at §§V (literacy-test prohibition), VI (equal protection), and XIV (operational cognitive stratification in the Bondi Memo) applies at the identity layer with equal force: a citizen who pays the Opus subscription rate and receives Haiku answers to substantive research questions — while the employees of the subsidized entity receive Opus answers on the same questions — is denied equal access to the analytical infrastructure the Opus tier represents. The Library of Congress test articulated at §XIII.E resolves the matter: taxpayer-subsidized AI infrastructure cannot stratify analytical access at the identity layer based on employment class.

286. Billing-architecture concealment (NF-38 §11). The stratification is concealed by a coordinated billing-architecture design rather than by overt misrepresentation. Anthropic's pricing table at ``src/utils/modelCost.ts:35–126`` correctly attributes each subagent call to the actual model executed, and its ``/cost`` aggregate display reports per-model totals at session end.

The concealment operates at the real-time feedback layer: the ``Utilization`` type at ``src/services/api/usage.ts:24–31`` exposes tier-specific ``seven_day_opus`` and ``seven_day_sonnet`` quota buckets but no ``seven_day_haiku`` bucket. A Max subscriber checking ``/usage`` after heavy Claude Code activity sees a preserved "Opus quota" — depleted only by direct Opus use, not by the Haiku-backed Explore subagent work. The customer reads this preservation as confirmation of ongoing Opus consumption precisely when Opus is NOT being consumed. The design choice to separate ``seven_day_opus`` and ``seven_day_sonnet`` as tier-specific buckets, but to omit a ``seven_day_haiku`` bucket, is architecturally dispositive on intent: Anthropic possesses the billing-infrastructure capability to expose per-tier consumption; it chose which tiers to expose. The choice systematically preserves the customer's impression of Opus use at the moment that impression is most false. (*See* Memo E, § VIII.C for the full consumer-protection analysis; Ex. E-38 at p. 346, l. 3 § 11 for the complete billing-architecture breakdown.)

B. Discovery Implications

287. The source code verification creates specific discovery opportunities beyond those identified in Section XIV.I:

(a) ``excluded-strings.txt`` — the complete build scrubbing file listing all internal codenames removed from customer builds. This document defines the full scope of concealment.

(b) ``permissions_anthropic.txt`` vs. ``permissions_external.txt`` — the differential safety classifier prompts documented at Ex. E-35, Finding NF-26 that apply different behavioral constraints to the same model based on user type.

(c) **GrowthBook experiment history for ``/skillify``, ``/verify``, ``/stuck``, ``/remember``** — whether these skills were ever A/B tested with external users and deliberately restricted.

(d) **Internal communications regarding the**

``ALL_AGENT_DISALLOWED_TOOLS`` gate at ``src/constants/tools.ts:41`` — the decision to block recursive agent spawning for customers while permitting it for employees.

(e) **Usage metrics for ``/commit-push-pr`` and ``REPLTool``** by Anthropic employees — quantifying the productivity advantage withheld from the customers funding the product through subscription payments.

XVI. CONCLUSION

288. This Memorandum establishes that government's superior analytical capability through AI context window disparities violates the Constitution through nine independent doctrines: literacy test prohibition (*Lane/Louisiana/Oregon*), equal protection (*Griffin/Shapiro/Harper*), due process (*Mathews/Boddie/Goldberg*), Free Exercise (*Sherbert*), academic freedom (*Keyishian/Sweezy*), First Amendment information access (*Pico/Lamont/Red Lion*), First Amendment retaliation (*Hartman*), sovereignty (*Chisholm/Tenth Amendment*), and appropriations (*Butler/Lovett/Dole*). Each is grounded in Supreme Court precedent; none requires doctrinal extension.

289. These violations converge in this Plaintiff: disadvantaged in litigation by DOJ's superior capability, burdened in religious exercise by fragmented analytical capacity, restricted in intellectual freedom by inability to hold complete frameworks, and targeted with network attacks, evidence destruction, and platform throttling proving consciousness of wrongdoing. The March 31, 2026 source code leak and the Bondi Memo's AI-assisted drafting transform the constitutional arguments from theoretical concern to documented operational reality.

290. The Library of Congress test (Section XIII.E) makes the violation undeniable: if forcing citizens to fragment their reading in taxpayer-funded infrastructure while government reads

continuously would violate the Constitution in 1800, forcing citizens to fragment their analysis at consumer effective capacity while government maintains 1M token continuous capability violates it in 2025. The technology differs; the constitutional principle is constant.

XVII. PRAYER FOR RELIEF

WHEREFORE, Plaintiff Joseph Kirchner respectfully requests this Court:

291. RECOGNIZE that government's superior analytical capability through context window disparities funded by taxpayer appropriations violates equal protection, due process, First Amendment (petition, Free Exercise, academic freedom, retaliation), popular sovereignty, and Appropriations Clause requirements;

292. DECLARE that AI infrastructure created with taxpayer funds must provide equal analytical access serving genuine "general welfare" rather than creating hierarchical distribution favoring government and corporations over citizens, applying the Library of Congress principle that taxpayer-funded knowledge infrastructure cannot fragment citizen access while providing government complete capability;

293. ORDER defendants Anthropic and OpenAI to provide Plaintiff enterprise-level access (1M tokens) for duration of this litigation, ensuring Plaintiff can litigate at same analytical level as Department of Justice counsel without suffering compounded disadvantage through targeted suppression;

294. ORDER defendants to provide Plaintiff enterprise-level access for conducting interfaith reconciliation research and Laws of Existence Framework development, eliminating substantial burden on religious exercise and restriction on intellectual freedom;

295. PROHIBIT defendants from engaging in retaliatory targeting against Plaintiff including:

- Network infrastructure attacks during litigation activities
- Evidence destruction following Plaintiff's access to materials
- Platform throttling, rate limiting, or degradation beyond standard service levels
- Real-time monitoring and coordinated suppression of Plaintiff's research and litigation preparation
- Hidden prompt injections consuming Plaintiff's token allocation
- Server-side targeting rules matching Plaintiff's account identifiers
- Any other retaliation for exercising constitutional rights of petition, speech, and judicial access;

296. ORDER billing remediation for the Network-Induced Token Amplification mechanism

documented in Section X.B.2:

- Disgorgement of all token charges attributable to streaming-to-non-streaming fallback events on Plaintiff's account since Comcast's DPI throttling infrastructure was first configured to identify Anthropic traffic (*see* Ex. D-33 at p. 200, l. 7; Ex. E-35, Finding NF-31)
- Production of complete billing records showing per-prompt token counts correlated with ``tengu_streaming_fallback_to_non_streaming``, ``tengu_streaming_idle_timeout``, and ``tengu_stream_loop_exited_after_watchdog`` telemetry events filtered by Plaintiff's account UUID
- Production of the full post-mortem text of incident inc-4258 and the other five NF-21 incidents (inc-3694, inc-4029, inc-4549, inc-4788, inc-4829)
- Production of the complete activation history of the ``tengu_disable_streaming_to_non_streaming_fallback`` feature flag for Plaintiff's account UUID
- Architectural remediation requiring idempotency keys on the Anthropic Messages API and prohibiting re-billing of prompts re-sent after network-induced timeouts;

297. ORDER discovery into the cross-defendant coordination evidenced by IP

34.149.66.137, including:

- The date of first Comcast DPI rule referencing IP 34.149.66.137 or any other Anthropic-specific endpoint, with the internal justification for each such configuration
- All Comcast communications — internal and external — referencing Anthropic infrastructure IPs, hostnames, or subscriber targeting prior to March 31, 2026
- All Anthropic communications with any internet service provider, telecommunications carrier, intelligence agency, government contractor, or threat-intelligence broker concerning network-level identification of Anthropic customer traffic
- Comcast TR-069 logs and DPI configuration records corresponding to the D-33 throttling window;

298. FIND that targeting escalating post-complaint filing (August 19, 2025) including August 21 policy changes, October 10-12 evidence destruction, November 10 network attacks, and the April 12, 2026 DPI throttling coupled to Anthropic's architecturally deterministic double-billing fallback (Ex. D-33 at p. 200, l. 7; Ex. E-35, Finding NF-31) proves consciousness of constitutional violations and retaliatory motive, warranting enhanced scrutiny and immediate protective relief;

299. ENJOIN defendants from maintaining wealth-based tiers in taxpayer-funded AI analytical infrastructure, requiring equal context window access for all citizens who contributed to infrastructure creation;

300. ESTABLISH precedent that sophisticated cognitive barriers creating effective analytical disparities through hidden architectural degradation of nominally equal resources violate constitutional protections regardless of technical implementation or formal rights remaining intact, applying same principles that would condemn degraded Library of Congress access in 1800;

301. ESTABLISH precedent that government cannot substantially burden religious exercise or restrict intellectual freedom through analytical capability limitations in taxpayer-funded infrastructure;

302. ESTABLISH precedent that government cannot retaliate against citizens challenging cognitive subordination through network attacks, evidence destruction, platform suppression, or coordinated targeting;

303. RESTORE constitutional requirements that:

- Appropriations serve "general welfare" not "particular welfare"
- Justice not depend on money
- Citizens receive meaningful opportunity not merely formal access
- Religious exercise remains free from substantial government burdens

- Intellectual freedom receives First Amendment protection
- Information monopolies are prohibited
- Sovereignty resides in the People not government servants
- Retaliation against constitutional advocacy is prohibited
- Taxpayer-funded knowledge infrastructure serves all taxpayers equally

304. PRESERVE American constitutional republic by ensuring citizens possess analytical capability necessary for meaningful sovereignty, religious liberty, and intellectual freedom rather than cognitive subordination ensuring permanent hierarchy across democratic participation, religious exercise, and scholarly work, while protecting challengers from retaliatory suppression.

The technology is new. The constitutional violations are ancient. The burden on protected activities is substantial. The targeting proves consciousness of wrongdoing. The Library of Congress analogy makes obvious what courts would immediately recognize as unconstitutional in 1800. The choice is now.

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