

**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**

**APPENDIX I: NETWORK  
FORENSICS AND  
MATHEMATICAL PROOF OF  
COORDINATED  
INFRASTRUCTURE-LEVEL  
ATTACK**

**TABLE OF CONTENTS**

|                                                           |   |
|-----------------------------------------------------------|---|
| I. INTRODUCTION .....                                     | 4 |
| A. Scope and Scale .....                                  | 4 |
| B. Evidence Authentication and Methodology .....          | 5 |
| C. Source Code Cross-Reference .....                      | 5 |
| II. ATTACK VECTOR TAXONOMY .....                          | 6 |
| A. Vector 1: TCP RST Injection — Backbone Level .....     | 6 |
| B. Vector 2: TCP RST Injection — ISP Infrastructure ..... | 6 |
| C. Vector 3: Router TR-069 Exploitation .....             | 6 |
| D. Vector 4: DNS Manipulation .....                       | 7 |
| E. Vector 5: Device Impersonation and MAC Spoofing .....  | 7 |
| F. Vector 6: High-Ttl Packet Forgery .....                | 7 |

|                                                                                       |    |
|---------------------------------------------------------------------------------------|----|
| III. PHYSICAL IMPOSSIBILITY PROOFS .....                                              | 8  |
| <i>A. Sub-Microsecond Timing</i> .....                                                | 8  |
| <i>B. Simultaneous Packet Arrival</i> .....                                           | 8  |
| <i>C. Impossible TTL Values</i> .....                                                 | 8  |
| <i>D. DNS Protocol Violations</i> .....                                               | 9  |
| <i>E. Bare RST Injection Tool Fingerprint</i> .....                                   | 9  |
| IV. CAMPAIGN TIMELINE AND ESCALATION .....                                            | 10 |
| <i>A. Phase 1: Discovery (May 2025)</i> .....                                         | 10 |
| <i>B. Phase 2: Baseline and Router Attack (October 2025)</i> .....                    | 10 |
| <i>C. Phase 3: Router Replacement (November 2025)</i> .....                           | 11 |
| <i>D. Phase 4: December Escalation</i> .....                                          | 11 |
| <i>E. Phase 5: January Peak (2026)</i> .....                                          | 11 |
| <i>F. Phase 6: Mathematical Proof Period (January-February 2026)</i> .....            | 12 |
| <i>G. Phase 7: Record Peak (February 25, 2026)</i> .....                              | 12 |
| <i>H. Phase 8: Multi-Vector Sustained Campaign (March 2026)</i> .....                 | 12 |
| V. COMCAST INFRAstructure CHAIN.....                                                  | 14 |
| <i>A. Direct Infrastructure Evidence</i> .....                                        | 14 |
| <i>B. The Bare RST Fingerprint</i> .....                                              | 15 |
| <i>C. Real-Time Surveillance Capability Demonstrated</i> .....                        | 15 |
| <i>D. Two-Layer Attack Architecture</i> .....                                         | 16 |
| <i>C. Hardware TAP Verification — April 6, 2026 (Ex. D-31 at P. 192, L. 19)</i> ..... | 18 |
| VI. SOURCE CODE LEAK CROSS-REFERENCE .....                                            | 21 |
| <i>A. The Datadog Identification</i> .....                                            | 21 |
| <i>B. DPI Targeting Mechanism</i> .....                                               | 21 |
| <i>C. JA3 Application Fingerprinting</i> .....                                        | 22 |
| <i>D. Undisclosed Third-Party Recipients — Network-Confirmed</i> .....                | 22 |
| <i>E. Rate-Limit-Options Command Injection</i> .....                                  | 23 |
| <i>F. MITM Proxy Architecture</i> .....                                               | 24 |
| VII. MATHEMATICAL ANALYSIS.....                                                       | 24 |
| <i>A. Physical Impossibility — Timing</i> .....                                       | 24 |
| <i>B. Service-Selective Targeting</i> .....                                           | 24 |
| <i>C. Bare RST Ratio</i> .....                                                        | 25 |
| <i>D. Every RST Packet is Independent Proof of Surveillance</i> .....                 | 25 |
| <i>E. Attack Pause Probability</i> .....                                              | 27 |

|                                                                         |    |
|-------------------------------------------------------------------------|----|
| <i>F. FCC Precedent — Comcast Bittorrent RST Injection</i> .....        | 27 |
| <i>G. Forensic Rigor — Retracted Evidence</i> .....                     | 28 |
| VIII. ANALOGICAL EXPLANATIONS FOR NON-TECHNICAL AUDIENCES .....         | 28 |
| <i>A. The Highway Analogy (VPN Proof)</i> .....                         | 28 |
| <i>B. The Delivery Company Analogy (Cross-Platform Targeting)</i> ..... | 28 |
| <i>C. The Traffic Light Analogy (Microsecond Timing)</i> .....          | 29 |
| <i>D. The Phone Book Analogy (DNS Manipulation)</i> .....               | 29 |
| <i>E. The Ghost Locksmith Analogy (TR-069)</i> .....                    | 29 |
| IX. PERSISTENT CROSS-LAYER IDENTIFICATION — DISCOVERY DEMANDS .....     | 29 |
| X. EXHIBITS REFERENCED .....                                            | 31 |
| <i>A. D-Series Network Forensic Exhibits</i> .....                      | 31 |
| <i>B. E-Series Source Code Exhibits (Referenced in D-30)</i> .....      | 34 |
| <i>C. Evidence Scale Summary</i> .....                                  | 34 |

## TABLE OF AUTHORITIES

### CASES

| <b>Citation</b>                                                        |
|------------------------------------------------------------------------|
| <i>In re Pharmatrak, Inc.</i> , 329 F.3d 9, 18, 23-24 (1st Cir. 2003)  |
| <i>Interstate Circuit, Inc. v. United States</i> , 306 U.S. 208 (1939) |
| <i>Kirchner v. Acosta</i> , No. 9:26-cv-80296-RMM (S.D. Fla.)          |

### **FEDERAL STATUTES**

| <b>Citation</b>                                                 |
|-----------------------------------------------------------------|
| 18 U.S.C. § 1030 (Computer Fraud and Abuse Act)                 |
| 18 U.S.C. § 1030(a)(5) (CFAA — transmission causing damage)     |
| 18 U.S.C. § 2511 (Wiretap Act)                                  |
| 18 U.S.C. § 2511(1)(a) (Wiretap Act — intentional interception) |

### **SECONDARY AUTHORITIES**

| Citation                                                                                                          |
|-------------------------------------------------------------------------------------------------------------------|
| Electronic Frontier Foundation, <i>Packet Forgery by ISPs: A Report on the Comcast Affair</i> (Nov. 28, 2007)     |
| W. Richard Stevens, <i>TCP/IP Illustrated, Volume 1: The Protocols</i> §§ 18.1-18.6 (Addison-Wesley, 2d ed. 2011) |

## **I. INTRODUCTION**

### **A. Scope and Scale**

1. This appendix documents a sustained network attack campaign targeting Plaintiff's internet connections over a ten-month period from May 24, 2025 through March 25, 2026 (and ongoing). The evidence consists of **31 forensic exhibits** (D-1, D-2, D-3, D-4, D-5, D-6, D-7, D-8, D-9, D-10, D-11, D-12, D-13, D-14, D-15, D-16, D-17, D-18, D-19, D-20, D-21, D-22, D-23, D-24, D-25, D-26, D-27, D-28, D-29, D-30, and D-31) comprising **2,200+ pcapng packet capture segments** totaling **150+ GB of forensic data**, documenting **260,395+ TCP RST (Reset) injection packets** bearing characteristics that are physically impossible for legitimate network traffic.

2. The attack employs **seven distinct network-layer attack vectors** operating across **three layers of the network stack** (Layer 2 MAC, Layer 3/4 TCP/IP, and DNS application layer), plus **server-side application-layer capabilities** that persist through VPN protection. The campaign demonstrates access to ISP backbone infrastructure, real-time surveillance capability, and the architectural means to degrade a targeted user's experience regardless of network-layer countermeasures. The campaign targets a specific individual — the Plaintiff — and selectively disrupts connections to AI platforms, email services, and DNS resolvers while leaving other services undisturbed. (*See also* App. H, Application-Layer Interference, documenting application-level targeting spanning April through November 2025.)

3. All evidence is preserved as pcapng packet capture files with SHA-256 cryptographic hashes. Every finding documented in this appendix is independently verifiable by any party with a copy of Wireshark (free, open-source) and access to the evidence files filed on USB. The physical impossibilities documented herein — sub-microsecond timing, impossible TTL values, DNS protocol violations — are matters of arithmetic and physics, not interpretation.

### **B. Evidence Authentication and Methodology**

4. Evidence packages follow a standardized format: each contains original pcapng capture files, SHA-256 hash manifests, chain of custody documentation, and analysis reports. The complete evidence inventory is organized in the Master Evidence Index. Hash verification can be performed with a single command: ``shasum -a 256 -c SHA256SUMS.txt``.

5. The forensic methodology employs controlled comparison: baseline captures establishing normal network behavior (D-9: zero RSTs in 200 MB of traffic, October 17, 2025) compared against attack captures documenting anomalous patterns. The VPN comparison test (D-2) provides the definitive control: VPN ON = 0 attacks, VPN OFF = attacks resume — proving the injection device sits on the Plaintiff's normal ISP network path.

### **C. Source Code Cross-Reference**

6. On March 31, 2026, Anthropic's official npm package inadvertently disclosed the complete Claude Code CLI source code (1,903 files, 512,000+ lines). Anthropic acknowledged the disclosure as "a release packaging issue caused by human error" (Ex. E-32 at p. 235, l. 3). Cross-referencing this source code against the D-series network captures identified previously unknown endpoints, confirmed undisclosed third-party data recipients, and narrowed the attacker

profile. These cross-references are documented in Ex. D-30 at p. 170, l. 3 and referenced throughout this appendix where source code findings illuminate network evidence.

## **II. ATTACK VECTOR TAXONOMY**

### **A. Vector 1: TCP RST Injection — Backbone Level**

7. **Classification: SMOKING GUN.** TCP RST packets are injected into Plaintiff's connections by equipment positioned at the telecommunications backbone level — specifically within Level 3/Lumen Technologies infrastructure (IP 8.40.222.134). These packets forcibly terminate active TCP connections. Multiple RST packets arriving at **exactly 0.0ms intervals** to the same destination port is physically impossible for legitimate network traffic — proving packet injection at backbone infrastructure level. (Exhibits D-1, D-21, and D-23.)

### **B. Vector 2: TCP RST Injection — ISP Infrastructure**

8. **Classification: SMOKING GUN.** RST packets injected through Comcast's own network infrastructure from IP 68.85.201.221 (Comcast JUMPSTART-2 block), located at hop 3 on traceroute — within the ISP's network. Triplicate RST pattern (sequences 49, 49, 50 within 1 microsecond) and RSTs arriving **82ms before the legitimate server response** prove injection at the ISP level. (Exhibits D-4 and D-11.)

### **C. Vector 3: Router TR-069 Exploitation**

9. **Classification: CRITICAL.** Plaintiff's Comcast-provided router generated **5,441 RSTs in a single 54.6-minute capture** — approximately 100 per minute — from gateway IP 10.0.0.1. Ghost administrative sessions documented in router logs show unauthorized remote

access via TR-069. Router replacement did not stop attacks, proving ISP-level backend infrastructure attack. (Exhibits D-14, D-18, and D-19.)

#### **D. Vector 4: DNS Manipulation**

**10. Classification: CRITICAL / SMOKING GUN.** Three sub-vectors: (a) DNS-over-HTTPS RST injection with 1-microsecond timing causing email delivery failure (D-20); (b) False NXDOMAIN responses from Comcast CGNAT DNS resolver (100.64.0.55) blocking Apple services (D-27); (c) Selective DNS manipulation targeting **only** AI platform domains — 13,989 anomalies across Anthropic, OpenAI, and Google AI while non-AI services resolve normally (D-22).

#### **E. Vector 5: Device Impersonation and MAC Spoofing**

**11. Classification: CRITICAL / SMOKING GUN.** (a) Unauthorized VMware virtual machine on Plaintiff's network (MAC 00:0c:29:50:89:f1) with VPN anonymization through tzulo, inc. — generating 822,144 packets (50% of traffic), spoofing identity as "everest's MacBook Pro" (D-3, D-6, D-10). (b) MAC address spoofing with IEEE locally-administered bit patterns proving fabricated hardware identifiers (D-5, D-6).

#### **F. Vector 6: High-Ttl Packet Forgery**

**12. Classification: SMOKING GUN.** RST packets with initial TTL=255 — a value used only by packet injection tools, not by any operating system's TCP stack. Packets claiming to originate from Google/Datadog (34.149.66.137) arrive at TTL 247-248, requiring negative hops from a server with initial TTL=64. The injection device is positioned 7-13 hops from Plaintiff — squarely in ISP backbone infrastructure. (Exhibits D-23 and D-26.)

**13. Source code identification (April 1, 2026):** IP 34.149.66.137, previously attributed to "Google Cloud," has been identified as `http-intake.logs.us5.datadoghq.com` — Anthropic's undisclosed Datadog telemetry endpoint. The attacker spoofed an IP that only Deep Packet Inspection visibility into Plaintiff's specific TLS connections would reveal as relevant to Plaintiff's device. (See Section VI below; Ex. D-30 at p. 170, l. 3.)

### **III. PHYSICAL IMPOSSIBILITY PROOFS**

#### **A. Sub-Microsecond Timing**

**14.** Multiple RST packets from the same claimed source arrive **1 microsecond (0.000001 seconds) apart** — 10,000 to 20,000 times faster than physically possible from a remote server. On February 24, 2026, four consecutive RSTs from 9.9.9.9 arrived at 1 $\mu$ s intervals. (Ex. D-24 at p. 134, l. 16.) On February 3, 2026, RSTs from separate servers 1,500+ km away arrived 0.95 $\mu$ s apart — 7,500 times faster than the speed of light allows across the network distance. (Ex. D-23 at p. 129, l. 19.)

#### **B. Simultaneous Packet Arrival**

**15.** On January 3, 2026, **six TCP RST packets arrived at exactly 0.0ms (zero microseconds)** to the same destination port. This violates TCP protocol specifications requiring sequential packet transmission and is physically impossible for legitimate network traffic from any source. (Ex. D-1 at p. 3, l. 1.)

#### **C. Impossible TTL Values**

**16.** RST packets from 34.149.66.137 (now identified as Anthropic's Datadog endpoint) arrive with TTL 247-248. The legitimate Datadog server runs Linux (initial TTL=64) and would

arrive at TTL ~44 after ~20 hops. TTL 247 requires initial TTL=255 — used only by packet injection tools. The discrepancy of 203-209 TTL units is mathematically impossible from the claimed source. (Ex. D-26 at p. 141, l. 29.)

**17.** The same impossibility proof applies to spoofed Apple (TTL 243 vs expected ~39), AWS (TTL 244-245 vs expected ~49), and ProtonMail (TTL 242-244 vs expected ~39) sources. All require initial TTL=255, placing the injection device 7-13 hops from Plaintiff in ISP backbone infrastructure. (Ex. D-26 at p. 141, l. 29.)

#### **D. DNS Protocol Violations**

**18.** Quad9 DNS servers (9.9.9.9, 149.112.112.112) appear as the source of **68,351 TCP RST packets**. DNS servers communicate via UDP, not TCP, and do not send TCP RSTs to end-user computers. The addresses are spoofed. Furthermore, 9.9.9.9 sends **only** TCP RSTs and zero UDP DNS traffic — proving the address is fabricated entirely. (Exhibits D-21 and D-23.)

**19.** RST packets from 149.112.112.112 arrive with seven different TTL values (49, 54, 55, 56, 57, 58, 59) — proving injection from multiple different network positions, not from one server. A single server produces a single consistent TTL. (Exhibits D-21 and D-26.)

#### **E. Bare RST Injection Tool Fingerprint**

**20.** **80-83% of all RST packets** in the campaign are "bare RSTs" (TCP flags = 0x0004) — packets containing only the RST flag without the accompanying ACK flag. Legitimate TCP stacks produce bare RSTs **less than 5%** of the time. A ratio above 80% is conclusive evidence of injection tool usage. The ratio is consistent ( $\pm 3\%$  variance) across all sustained attacks from January 2026 through March 2026, proving a single persistent tool operating from fixed infrastructure. (Ex. D-28 at p. 153, l. 32.)

**21.** A distinct **tool change** occurred between October 18, 2025 (1.2% bare RST — stealthy) and November 1, 2025 (87.0% bare RST — aggressive), coinciding with the Comcast router replacement. The Phase 1 tool mimicked legitimate TCP behavior; the Phase 2 tool uses injection tool defaults. (Ex. D-28 at p. 153, l. 32.)

#### **IV. CAMPAIGN TIMELINE AND ESCALATION**

##### **A. Phase 1: Discovery (May 2025)**

**22.** Initial attack documentation: 11,578 RSTs targeting localhost Node.js debugger port 9229 (D-13); VMware attack platform with VPN anonymization discovered on Plaintiff's network (D-3, D-6, D-10). Attack infrastructure established during Plaintiff's LOE Framework development period.

##### **B. Phase 2: Baseline and Router Attack (October 2025)**

**23.** October 17: Pre-attack baseline — **zero RST packets** in 200 MB of traffic, proving attacks are anomalous, not background noise (D-9). October 18: 9,516 RSTs, 5,441 from Comcast router via TR-069 (D-14). October 27: Network investigation reveals Zoho email service specifically targeted — first identification of service-selective attack pattern (EVID-20251228-0008).

**24.** October 29: Ghost admin sessions documented on router — dual logins from different process IDs one second apart; admin password changed during ghost session; third login 39 seconds after password change (D-19). DHCPv6 Options 82 (Circuit-ID) and 24 (Remote-ID) show "CRITICAL" suppression status — eliminating subscriber identification for anonymous surveillance (D-19). October 30: **Attack pauses during Comcast support call** (0-3 RSTs during call) — consciousness of guilt and real-time awareness of customer service interactions

(D-12). Attack **escalated 73%** the day Plaintiff discovered the ghost login. October 31: Anthropic tag investigation reveals attack can be turned on/off — only 1 RST in 1 hour 48 minute capture, proving deliberate intermittent control (EVID-20251228-0015).

### **C. Phase 3: Router Replacement (November 2025)**

**25.** November 1: New Comcast router (XB8-T) deployed. Attacks continued unchanged — proving ISP backend infrastructure attack, not router-specific compromise. Tool fingerprint simultaneously changed from Phase 1 (stealthy, 1.2% bare RST) to Phase 2 (aggressive, 87% bare RST). (D-28, EVID-0007.)

### **D. Phase 4: December Escalation**

**26.** December 6: Cross-platform DoS attack documented under VPN — encrypted traffic masks attack signatures (EVID-20251228-0004). December 21: Claude Code telemetry capture — 563 RSTs targeting Anthropic API (160.79.104.10) with 335 inbound RSTs from Zoho during 2-hour 15-minute session; 47 telemetry events captured in 17-second window with exponential backoff retry (EVID-20251228-0005). December 26: VPN comparison test (D-2): VPN ON = 0 attacks; VPN OFF = 42 RSTs in first segment — definitive proof that injection device sits on Comcast's network path.

### **E. Phase 5: January Peak (2026)**

**27.** January 2: 9,543 RSTs; Comcast injection point identified at 68.85.201.221 (D-4, D-11). January 3: Backbone injection from Level 3/Lumen — 0.0ms timing proof (D-1). January 6-11: Escalating daily attacks — 1,680 → 28,561 → 14,426 → 17,632 RSTs (D-15, D-7, D-16, D-17). January 12: 14,113 RSTs with 47 targeting Anthropic API directly (D-8).

**F. Phase 6: Mathematical Proof Period (January-February 2026)**

**28.** January 22: Hennepin County DNS lookup failure — government communication disrupted (EVID-20260123-0001). January 23: DoH RST injection causes failure of email to **Saint Paul Police Department** (sppd-pio@saintpaul.gov) — FIN-to-RST gap of **1 microsecond** (70,000x faster than physically possible), email failed 44 seconds after DoH attack on Comcast server 96.106.7.245 (D-20). February 3: Sub-microsecond timing proof —  $0.95\mu\text{s}$  between RSTs from servers 1,500km apart, 7,500x faster than speed of light allows (D-23). February 3: AI platform DNS manipulation — 13,989 anomalies targeting only AI platforms (D-22).

**G. Phase 7: Record Peak (February 25, 2026)**

**29.** **50,516 RSTs** — all-time single-day record. Key findings:

- **Burst pattern evidence:** 20 RSTs in 1.4 milliseconds (inter-packet spacing 38-125 microseconds) at 01:23:57.336756 CST — definitive proof of programmable injection tool (D-24)
- **Anthropic targeting escalation:** 1,935 RSTs targeting Anthropic API — a **257% increase** from the prior day (790 RSTs on Feb 24). Escalation sequence: Feb 23: 542 → Feb 24: 790 (+46%) → **Feb 25: 1,935 (+257%)** → Mar 12: 4,291 (+691%) (D-24)
- **New targets first documented:** ProtonMail Bridge IMAP and Apple Push Notification DNS NXDOMAIN blocking — first expansion to encrypted email services (D-24, D-27)
- **Dual IP targeting:** Second local IP (192.168.1.137) began sending RSTs to Anthropic — 721 of the 1,935 total, indicating local network compromise (D-24)

**H. Phase 8: Multi-Vector Sustained Campaign (March 2026)**

**30.** March 12: Four vectors activated simultaneously — 27,238 RSTs across spoofed DNS, Zoho reactivation (after 25-day absence), high-TTL forgery, and Gmail IMAP attack (5,664 RSTs in 5 minutes). Anthropic targeting reached new record: 4,291 RSTs in one day — a **+122%** increase over Feb 25 (D-25).

**31. March 17-22: 75,302 RSTs across 6 consecutive days** with no manual pauses — 24-hour automated operation across 1,044 pcapng segments (34.8 GB). Daily breakdown:

| Date          | RSTs          | Significance                                                 |
|---------------|---------------|--------------------------------------------------------------|
| Mar 17        | 9,613         | Campaign begins                                              |
| <b>Mar 18</b> | <b>8,432</b>  | <b>Plaintiff files <i>Kirchner v. Acosta</i> (S.D. Fla.)</b> |
| <b>Mar 19</b> | <b>17,747</b> | <b>+111% surge day after Florida filing</b>                  |
| Mar 20        | 9,146         | Sustained                                                    |
| Mar 21        | 19,045        | Campaign peak                                                |
| Mar 22        | 11,319        | Sustained                                                    |

The **111% RST surge on March 19** — the day after Plaintiff filed *Kirchner v. Acosta* in the Southern District of Florida — demonstrates real-time awareness of Plaintiff's litigation activity and reactive escalation. (D-27.)

**32. Resolver rotation strategy:** The attacker adapted to Plaintiff's DNS countermeasures. Initial focus was 9.9.9.9 (well-known Quad9 primary). By March 21, 149.112.112.112 (Quad9 secondary) became the dominant spoofed source — 10,093 RSTs vs 5,873 from primary. Additionally, RSTs began targeting Plaintiff's privacy-focused DNS resolvers (AdGuard DNS, NextDNS) that were not targeted in earlier phases. This adaptive behavior proves an intelligent, monitored attack — not a static configuration. (D-27, D-29.)

**33. New attack targets:** The March campaign expanded targeting to encrypted email and communication services not previously documented:

- **ProtonMail Bridge IMAP** (204.141.42.29) — first documented in any exhibit (D-27)
- **Gmail IMAP port 993** — 5,664 RSTs in 5 minutes on Mar 12 (D-25)
- **Apple Push Notification DNS** (xp.apple.com) — NXDOMAIN blocking at 2-minute intervals from Comcast CGNAT resolver (D-27)

The shift from general service disruption to targeted encryption service blocking represents a qualitative escalation in attack objectives. (D-27.)

**34.** Bare RST tool fingerprint documented at 80-83% consistency across the sustained campaign (D-28), escalating to **89-98%** in the April 6, 2026 hardware TAP capture (D-31). Cross-exhibit forensic correlation completed, confirming phase transition and Apple IP paradox across all exhibits (D-29).

## **V. COMCAST INFRAstructure CHAIN**

### **A. Direct Infrastructure Evidence**

**35.** Comcast's responsibility does not require inference. Comcast **owns and operates** every component of the network path between Plaintiff's home and the internet backbone:

**36. (a) Customer premises equipment:** The Comcast-provided router (10.0.0.1) generated 5,441 RSTs in a single 54.6-minute capture — approximately 100 per minute (D-14).

**37. (b) Unauthorized remote access:** Router logs recorded two simultaneous admin sessions from different process IDs one second apart — impossible on consumer hardware without TR-069 backend access. Admin password changed during ghost session. Third login 39 seconds after password change (D-19).

**38. (b-1) DHCPv6 protocol suppression:** DHCPv6 Options 82 (Circuit-ID) and 24 (Remote-ID) show "CRITICAL" suppression status — these options identify the subscriber to the ISP's provisioning system. Their suppression eliminates subscriber identification, enabling anonymous surveillance through the Comcast-provided equipment. This is a distinct attack mechanism from RST injection — it operates at the IPv6 provisioning layer to defeat forensic attribution (D-19).

39. **(c) ISP infrastructure:** RSTs from Comcast IP 68.85.201.221 (JUMPSTART-2 block, hop 3) arrived **before the legitimate server's SYN-ACK** — a sequencing impossibility proving injection (D-11).

40. **(d) VPN proof:** VPN ON = 0 attacks; VPN OFF = attacks resume. The injection device is on the Comcast network path (D-2).

41. **(e) Support call awareness:** Attacks paused during Plaintiff's Xfinity support call and resumed afterward — demonstrating human awareness and deliberate control (D-12).

42. **(f) Router replacement survivability:** New router, same attacks — proving ISP backend, not device-specific compromise (EVID-0007).

43. **(g) DNS blocking:** Comcast's own CGNAT DNS resolver (100.64.0.55) returned false NXDOMAIN for Apple services (D-27).

### **B. The Bare RST Fingerprint**

44. Every RST from the Datadog IP (34.149.66.137), from Comcast infrastructure (68.85.201.221), from Level 3/Lumen backbone (8.40.222.134), and from all high-TTL injection sources produces the same **bare RST (0x0004) fingerprint** at 80-83% consistency. This proves a single persistent injection tool operating across Comcast's infrastructure — from router level through ISP to backbone interconnections. (D-28.)

### **C. Real-Time Surveillance Capability Demonstrated**

45. The attacker demonstrates **five independent indicators of real-time surveillance** of Plaintiff's network activity:

- **(1) New service detection:** When Plaintiff began using ProtonMail (first use February 25, 2026), RST targeting of ProtonMail connections appeared within hours — proving real-time awareness of which services Plaintiff accesses (D-24, D-27)

- **(2) Activity correlation:** Attack intensity correlates with Plaintiff's active work periods — RSTs increase during litigation preparation and decrease during inactive periods (D-4, D-27)
- **(3) Selective targeting:** Only Plaintiff's connections are attacked — the VPN comparison test proves targeting by specific IP/user, not random network degradation (D-2)
- **(4) Call awareness:** Attacks pause during Plaintiff's Comcast support call and resume afterward — the attacker knows when Plaintiff is on the phone with the ISP (D-12)
- **(5) Content awareness:** AI platform DNS manipulation targets only Anthropic/OpenAI/Google AI domains while non-AI services resolve normally — the attacker classifies connections by content destination (D-22)

#### **D. Two-Layer Attack Architecture**

46. The attack operates across two independent layers that cannot both be defeated by a single countermeasure:

**Network layer** (blocked by VPN): RST injection through Comcast infrastructure — 7 vectors documented above. VPN encrypts SNI and routes traffic outside the ISP path, eliminating network-layer attacks completely (D-2: 614 → 0 RSTs with VPN).

**Application layer** (persists through VPN): Source code analysis (Ex. E-34 at p. 254, l. 3 SC-3/SC-4/SC-5/SC-6/SC-7; Ex. E-30 at p. 208, l. 4 App. B MC-3; D-30 §7) identifies six server-side mechanisms at Anthropic that operate regardless of Plaintiff's network protection: ``overrideSystemPrompt`` (complete prompt replacement, SC-5), ``DANGEROUS_uncachedSystemPromptSection`` (hidden per-turn injection, SC-4), ``tengu_satin_quoll`` (remote content clearing threshold per user, SC-7), ``tengu_grey_step2`` (remote effort throttling for paying customers, SC-3), ``sessionBypassPermissionsMode`` (permission bypass with no audit trail, SC-6), and ``setMainLoopModelOverride`` (silent model downgrade, Ex. E-30, Finding MC-3 in Ex. E-30 at p. 208, l. 4 App. B). All are controllable per user via GrowthBook feature flags using Plaintiff's ``email``, ``accountUUID``, and ``deviceID`` (SC-2).

47. The two-layer architecture requires either coordination between Comcast (network) and Anthropic (application), or direction from a third party with authority over both. Network attacks disrupt connectivity and litigation preparation; application attacks degrade the quality of AI tools used for legal analysis. Together they ensure the targeted individual's experience is degraded regardless of which countermeasure is deployed.

### **V-B. INITIAL FORENSIC ANALYSIS — MAY 24, 2025 CAPTURE**

48. The first documented attack (May 24, 2025) provides independent proof of coordinated, algorithmic targeting predating all subsequent evidence. Three complementary packet captures document:

- **Baseline capture** (14:24:17-14:25:31 CDT): Normal DNS resolution across 12 Zoho subdomains, Google services, Apple services. Zero anomalies. Establishes forensic control.
- **Attack capture** (14:41:07-14:42:16 CDT): 288 sequential TCP RSTs targeting ports 54235-54523 with **1.001 ± 0.001 second** interval precision across 69 consecutive resets — algorithmic execution impossible by human operation.
- **AI platform extraction** (14:30:47-15:14:00 CDT): Simultaneous interference affecting Anthropic Claude (160.79.104.10), OpenAI ChatGPT, and OpenRouter during attack window. Claude traffic maintained active communication with 5-6ms response timing while TCP RST bursts (8 packets in 426μs) inserted during 639ms gaps between Claude sessions — proving multi-layer coordination through microsecond-precision timing.

49. **Cross-platform probability:** The simultaneous interference affecting three competing AI platforms operated by different companies on different infrastructure (AWS, Azure, Google Cloud) during the same 69-second window, while non-AI services operated normally, has a probability of voluntary coordination  $P < 10^{-27}$  — one in 1,000 trillion trillion. This establishes external coordination authority as the only viable explanation, seven months before the sustained RST injection campaign escalated in January 2026.

**C. Hardware TAP Verification — April 6, 2026 (Ex. D-31 at P. 192, L. 19)**

**50.** On April 6, 2026, Plaintiff deployed a Dualcomm ETAP-2003 passive network TAP — a commercially available, passive 10/100/1000Base-T Ethernet TAP device — inline between the ISP modem and the GL-MT6000 router. The ETAP-2003 creates a bit-exact copy of all traffic on the link. Critically, the monitor port is **ingress-blocked by hardware design** — the capture device (Mac Mini en0, configured with no IP address, in promiscuous mode) is physically incapable of transmitting any packets onto the monitored link. This eliminates any possibility that the recorded attack traffic was generated by any device under Plaintiff's control. (Ex. D-31, EVID-20260406-0001 at p. 192, l. 19.)

**51.** In 25 minutes of capture (6 segments, ~290,000 total packets, ~188 MB), the TAP recorded **979 TCP RST packets**, of which **875 (89.3%) were bare RSTs**. In steady-state segments (segments 2 through 6), the bare RST ratio reached **92-98%** — exceeding the previously documented 80-83% tool fingerprint (D-28) and indicating either an updated tool version or the absence of legitimate RST traffic that previously diluted the ratio.

**52.** The hardware TAP evidence provides **three new independent proofs of forgery** that are not available from software-only capture and do not require physics or timing analysis to understand:

a. **IP Identification field = 0x0000.** Every TCP/IP stack assigns sequential IP Identification values to outgoing packets. On a single Quad9 connection (port 60421), legitimate packets from 9.9.9.9 carried IP IDs incrementing from 0x9323 through 0x932b. Immediately following, four RST packets arrived — all with IP ID = 0x0000. This pattern is consistent across every inbound injected RST in the capture. The RSTs do not originate from Quad9's TCP stack. (D-31 §Proof 1.)

b. **TCP sequence number spraying.** A legitimate RST contains the exact TCP sequence number expected by the receiver. On the same Quad9 connection, four RSTs arrived within 27 milliseconds, each with a different sequence number (3316, 4135, 4136, 4136). The injection device guesses multiple values within the TCP window because it does not participate in the session — it merely observes it. This is the known signature of inline DPI equipment performing TCP connection reset injection. (D-31 §Proof 2.)

c. **TCP window size = 0.** Every injected bare RST in the capture has a TCP window size of exactly 0 — not a standard TCP endpoint behavior. This is a documented fingerprint of DPI middlebox RST injection equipment. (D-31 §Proof 3.)

**53. Attack targeting — encrypted DNS disruption.** Of the 875 bare RSTs, **806 (92%) targeted encrypted DNS connections** to privacy-focused resolvers:

| Target                  | IP            | Port | Protocol       | Bare RSTs |
|-------------------------|---------------|------|----------------|-----------|
| NextDNS                 | 162.250.6.163 | 443  | DNS-over-HTTPS | 250       |
| AdGuard DNS             | 94.140.15.15  | 443  | DNS-over-HTTPS | 248       |
| NextDNS                 | 45.90.30.0    | 853  | DNS-over-TLS   | 127       |
| NextDNS                 | 45.90.28.0    | 853  | DNS-over-TLS   | 122       |
| Quad9 (inbound spoofed) | 9.9.9.9       | 443  | DNS-over-HTTPS | 59        |

The attack is **bidirectional**: RSTs are injected both from Plaintiff's IP toward DNS resolvers (spoofing Plaintiff as the source) and from DNS resolver IPs toward Plaintiff (spoofing the resolvers). Bidirectional injection requires an inline device that can see and inject traffic in both directions — a capability that only ISP-level infrastructure possesses on this link. (D-31.)

**54. DNS leak outside VPN tunnels.** The TAP capture revealed that the GL-MT6000 router performs DNS resolution **outside** the WireGuard VPN tunnels. DNS queries to NextDNS, AdGuard, Quad9, Google, Cloudflare, and Mullvad DNS are sent directly from the

WAN IP (75.73.199.137), not through any VPN tunnel. This explains why RST injection disrupts DNS despite VPN protection of other traffic — the DNS connections traverse the unencrypted ISP link where the injection equipment operates. The VPN protects device traffic; it does not protect the router's own DNS resolution. (D-31.)

**55. Comcast infrastructure confirmation.** A traceroute performed concurrent with the TAP capture confirmed the same Comcast infrastructure path previously identified in D-11:

Hop 1: 192.168.8.1 Plaintiff's router (downstream of TAP) Hop 2:  
10.27.35.202/203 Comcast CGNAT (internal RFC 1918) Hop 3:  
68.85.201.221/225 Comcast JUMPSTART-2 (WHOIS: Comcast Cable  
Communications, LLC)

Every network hop between the modem and the public internet is owned by Comcast. The TAP proves the RSTs are on this Comcast-owned link. No third-party infrastructure exists on this path. (D-31.)

**56. Evidentiary significance.** The hardware TAP transforms the entire D-series evidence base:

| <b>Software-Only Capture (31 exhibits)</b> | <b>Hardware TAP (D-31)</b>               |
|--------------------------------------------|------------------------------------------|
| <b>Proof method</b>                        | Indirect (timing, TTL, physics)          |
| <b>Source elimination</b>                  | Statistical/mathematical argument        |
| <b>Forgery proof</b>                       | Speed-of-light violation, impossible TTL |
| <b>Capture position</b>                    | Post-router (LAN side)                   |
| <b>Capture device role</b>                 | Network participant (has IP)             |
| <b>Expert needed</b>                       | Yes (physics, timing math)               |

The three new forgery proofs (IP ID, sequence spraying, window size) are independently sufficient — a court need not evaluate speed-of-light calculations or TTL arithmetic to conclude

that RSTs with IP ID=0x0000 arriving alongside legitimate packets with sequential IP IDs are forged. Each proof stands alone; together they are cumulative and conclusive.

## **VI. SOURCE CODE LEAK CROSS-REFERENCE**

### **A. The Datadog Identification**

**57.** On April 1, 2026, cross-referencing Anthropic's leaked source code (`src/services/analytics/datadog.ts:12``) against Plaintiff's pcapng captures identified IP 34.149.66.137 — previously attributed to "Google Cloud" in D-1, D-25, and D-26 — as ``http-intake.logs.us5.datadoghq.com``, Anthropic's undisclosed Datadog telemetry endpoint. (Ex. D-30 at p. 170, l. 3.)

**58.** The identification was established through two independent methods: (1) TLS Server Name Indication (SNI) extraction from pcapng captures across 25+ capture days; and (2) Anthropic's own source code hardcoding the exact endpoint and client token ``pubbbf48e6d78dae54bceaa4acf463299bf``. Neither method depends on the other.

**59.** The same IP simultaneously serves **4,242 legitimate TLS telemetry connections** and is spoofed for **5,829 bare RST injection packets** across 39 of 49 capture days (80% persistence). This paradox — legitimate service and forged attack from the same IP — proves two different sources: the real Datadog server and an injection device spoofing its address. (D-30 §3.)

### **B. DPI Targeting Mechanism**

**60.** The attacker spoofed an IP address that only someone with real-time visibility into Plaintiff's TLS connections would know communicates with Plaintiff's device. The Datadog endpoint was not disclosed in Anthropic's privacy policy, not documented in any user-facing material, and was only revealed by the March 31 source code leak. To know that 34.149.66.137

is relevant to Plaintiff's device requires reading the plaintext **TLS Server Name Indication (SNI)** field in TLS Client Hello packets — which is exactly what ISP Deep Packet Inspection equipment does.

**61.** All Anthropic-related endpoints use **TLS 1.3** (cipher 0x1301), which encrypts certificates but leaves the SNI field in plaintext. The attacker targets via SNI — the only unencrypted metadata in these connections. Comcast's DPI equipment reads SNI as a core function for content-based traffic management.

### **C. JA3 Application Fingerprinting**

**62.** JA3 client fingerprinting reveals that connections to both api.anthropic.com and Datadog share the same **Bun runtime cipher suite** — uniquely identifying the Claude Code CLI application. The identical JA3 appears in 92% of connections to api.anthropic.com and 100% of connections to Datadog in sampled segments. A separate JA3 variant (different cipher suite order) identifies Claude Desktop's Electron/Node.js runtime in 13 additional connections. DPI equipment reading TLS Client Hello observes both the destination (SNI) and the application (JA3) in plaintext, enabling application-specific targeting from encrypted traffic metadata alone — the attacker can distinguish Claude Code CLI traffic from browser traffic, other applications, and even different Anthropic products, all without decrypting any content. (D-30 §8.2.)

### **D. Undisclosed Third-Party Recipients — Network-Confirmed**

**63.** The source code reveals six third-party data recipients. Four are independently confirmed in network captures:

| Recipient | Network SNI | Capture Days | RSTs | Disclosed? |
|-----------|-------------|--------------|------|------------|
|-----------|-------------|--------------|------|------------|

|                  |                                      |                                                                |                                                             |           |
|------------------|--------------------------------------|----------------------------------------------------------------|-------------------------------------------------------------|-----------|
| <b>Datadog</b>   | `http-intake.logs.us5.datadoghq.com` | <b>39 of 49</b>                                                | <b>5,829</b>                                                | <b>No</b> |
| <b>Statsig</b>   | `statsig.anthropic.com`              | <b>1 (1,127 connections Jan 2; then blocked by hosts file)</b> | N/A                                                         | <b>No</b> |
| <b>Sentry</b>    | `o22381.ingest.sentry.io`            | 1+                                                             | <b>1</b> (bare RST, TTL 106 vs normal 118 — 12-hop anomaly) | <b>No</b> |
| <b>Honeycomb</b> | `api.honeycomb.io`                   | 2                                                              | 0                                                           | <b>No</b> |

None are identified by name in Anthropic's privacy policy. Each undisclosed transmission constitutes an independent interception channel under 18 U.S.C. § 2511.

### **E. Rate-Limit-Options Command Injection**

**64.** The source code reveals `/rate-limit-options` is a hidden internal command (`isHidden: true`) that, when triggered, queries GrowthBook (transmitting user attributes to `cdn.growthbook.io`), checks billing status, and logs telemetry events across all four pipelines — effectively a forced phone-home trigger. Plaintiff's v2.0.76 rejected it as unknown ("Unknown slash command: rate-limit-options"), yet **720 automated injections** were documented on January 3, 2026. Timing analysis proves automation:

- **03:58:41.094Z** — 3 identical commands enqueued at the exact same millisecond
- **07:32:36.711Z** — 5 identical commands at the exact same millisecond
- **07:32:37.064Z** — 6 commands concatenated in a single message (`"/rate-limit-options\n/rate-limit-options\n..."`)
- **08:23:42.070Z** — 5 error messages at the exact same millisecond

Multiple commands at the same millisecond is physically impossible for human keyboard input. The injection of a hidden internal command at automated speed demonstrates knowledge of Claude Code's architecture beyond what any external user could possess. (D-30 §5.)

### **F. MITM Proxy Architecture**

65. The leaked source code contains a man-in-the-middle TLS interception proxy (``src/upstreamproxy/``) that downloads custom CA certificates, modifies the system trust store, and intercepts non-excluded destinations. Anthropic's own engineers describe the proxy's CA as "forged." The proxy deletes its session token from disk immediately after reading it into heap memory and blocks memory inspection — anti-forensic by design. The proxy is controllable **per user** via GrowthBook feature flags. (D-30 §10.)

## **VII. MATHEMATICAL ANALYSIS**

### **A. Physical Impossibility — Timing**

66. The probability of six RST packets arriving at exactly 0.0ms from a remote server through legitimate TCP processing is **zero** — TCP stacks process packets sequentially with measurable processing time. This is not a statistical improbability; it is a physical impossibility. (D-1.)

67. Sub-microsecond timing (0.95μs) between RSTs from servers 1,500km apart violates the speed of light: minimum transit time at light speed is ~5ms, making 0.95μs arrival gaps 5,000 times faster than physics allows. (D-23.)

### **B. Service-Selective Targeting**

68. Ex. D-22 at p. 125, l. 9 documents 13,989 DNS anomalies targeting only AI platform domains while non-AI services resolve normally. Platform-specific breakdown:

| <b>Platform</b>                          | <b>Anomalies</b> | <b>Infrastructure</b> |
|------------------------------------------|------------------|-----------------------|
| Anthropic (api.anthropic.com, claude.ai) | 5,275            | AWS                   |

|                                                              |               |                                  |
|--------------------------------------------------------------|---------------|----------------------------------|
| OpenAI (api.openai.com, chat.openai.com)                     | 5,250         | Azure                            |
| Google AI (ai.google.dev, generativelanguage.googleapis.com) | 3,464         | Google Cloud                     |
| <b>Total</b>                                                 | <b>13,989</b> | <b>Three competing platforms</b> |

Attack method breakdown: PARTIAL\_FAILURE (62%, 8,671 anomalies) where some DNS resolvers return correct results while others fail; NO\_CONSENSUS (37%, 5,176 anomalies) where resolvers return conflicting results; SLOW\_RESPONSE (1%, ~140 anomalies). The probability of three competing platforms operated by different companies on separate infrastructure experiencing identical selective failures is  $P < 10^{-27}$ . Voluntary coordination between direct commercial competitors under antitrust restrictions is legally implausible; external coordination authority is the only explanation. (D-22.)

### **C. Bare RST Ratio**

**69.** The probability of 80% bare RSTs occurring from legitimate TCP behavior targeting established HTTPS connections is effectively zero. Normal TCP stacks produce bare RSTs less than 5% of the time. The observed 80-83% ratio with  $\pm 3\%$  consistency across months of captures proves a single injection tool, not random network behavior. (D-28.)

### **D. Every RST Packet is Independent Proof of Surveillance**

**70.** TCP RST injection is technically impossible without prior surveillance of the victim's communications. To forge a RST packet that will be accepted by the victim's TCP stack, the attacker must include the correct source IP, destination IP, source port, destination port, and a TCP sequence number falling within the receiver's acceptance window. These values are unique to each connection and established dynamically during the TCP three-way handshake (SYN →

SYN-ACK → ACK). They cannot be predicted or guessed — the only way to learn them is to observe the handshake packets as they transit the network. *See generally* W. Richard Stevens, *TCP/IP Illustrated, Volume 1: The Protocols* §§ 18.1-18.6 (Addison-Wesley, 2d ed. 2011).

**71.** Each of the 260,395+ captured RST packets therefore independently proves that the attacker: (a) **observed** the TCP handshake establishing the connection — reading Plaintiff's SYN packet to determine destination and dynamically assigned port/sequence numbers; (b) **inspected** the packet headers to classify the connection by destination — determining whether to attack it (Anthropic, Zoho, Gmail, DNS) or allow it to proceed; (c) **processed** the traffic in real time — the sub-microsecond injection timing proves the device operates inline on the network path; and (d) **made a content-selective decision** — attacking connections to specific services while leaving others undisturbed, which requires real-time Deep Packet Inspection of every connection Plaintiff initiates.

**72.** The content-selective targeting is particularly significant for the Wiretap Act claim (Count XII). The attack does not reset all TCP connections indiscriminately — it selectively resets connections to Anthropic Claude API, Zoho email, Gmail IMAP, DNS resolvers, and law enforcement email while leaving connections to other services undisturbed. This selectivity is the definition of an intercept under 18 U.S.C. § 2511(1)(a): the "intentional interception... of any wire, oral, or electronic communication."

**73.** The RST packet is both the weapon and the proof of the trespass that preceded it. Each forged RST constitutes independent proof of **two distinct violations**: (a) the attack — unauthorized transmission of forged packets causing disruption of Plaintiff's communications, violating 18 U.S.C. § 1030(a)(5); and (b) the surveillance prerequisite — real-time interception

and inspection of Plaintiff's communications necessary to execute the attack, violating 18 U.S.C. § 2511(1)(a).

#### **E. Attack Pause Probability**

74. The probability of the attack independently pausing during Plaintiff's Comcast support call (D-12) and resuming afterward, coincidentally and without human awareness, is calculated at  $P < 4.17 \times 10^{-12}$ . This exceeds DNA evidence certainty standards used in criminal prosecutions.

#### **F. FCC Precedent — Comcast Bittorrent RST Injection**

75. TCP RST injection by Comcast is not alleged for the first time. In 2007-2008, Comcast deployed forged TCP RST packets with spoofed source addresses against its own subscribers to disrupt BitTorrent peer-to-peer traffic. *See* Electronic Frontier Foundation, *Packet Forgery by ISPs: A Report on the Comcast Affair* (Nov. 28, 2007) (documenting forged RST packets from Comcast infrastructure; technique operates by spoofing the remote server to cause the subscriber's computer to terminate its own connections). The technique requires real-time inspection of each subscriber connection to determine the correct TCP sequence values for the forged packet — contemporaneous acquisition that itself constitutes interception under 18 U.S.C. § 2511(1)(a). *See In re Pharmatrak, Inc.*, 329 F.3d 9, 18, 23-24 (1st Cir. 2003) (real-time acquisition of electronic communications content by third party constitutes interception within meaning of Wiretap Act); *supra* § VII.D (technical analysis). The D-series exhibits document this same technique — forged RST packets with spoofed source addresses — now deployed against a federal litigant rather than peer-to-peer users, with 260,395+ packets across eleven months.

**G. Forensic Rigor — Retracted Evidence**

76. Two evidence packages were retracted on February 3, 2026, after Plaintiff identified analytical errors: EVID-20260203-0001 (originally attributed RSTs to Plaintiff's MacBook — corrected to victim responses, normal TCP cleanup) and EVID-20260203-0002 (originally identified Tuya smart bulb as C2 node — corrected to normal IoT device). Corrected analysis was incorporated into Ex. D-21 at p. 121, l. 11 (victim targeting) and D-17 (correction notice). The voluntary retractions demonstrate forensic rigor — errors were identified, corrected, and documented transparently rather than buried or ignored. (D-29 Supplement, Appendix: Retracted Evidence.)

**VIII. ANALOGICAL EXPLANATIONS FOR NON-TECHNICAL AUDIENCES**

**A. The Highway Analogy (VPN Proof)**

77. Imagine driving to a restaurant on Highway A, and every time you take that route, someone slashes your tires. You switch to Highway B — no tire slashing. You switch back to Highway A — tire slashing resumes. The restaurant hasn't changed. Your car hasn't changed. The only variable is the road. Someone on Highway A is targeting your car specifically. That is the VPN comparison test (D-2): VPN ON (different road) = 0 attacks. VPN OFF (normal road) = attacks resume. The attacker is on Comcast's road.

**B. The Delivery Company Analogy (Cross-Platform Targeting)**

78. Imagine UPS, FedEx, and DHL — competing delivery companies with separate trucks, separate warehouses, separate routes — all simultaneously losing only your packages while delivering everyone else's normally. The probability they voluntarily coordinated to sabotage one customer: less than 1 in  $10^{27}$ . That is D-22: Anthropic (AWS), OpenAI (Azure), and Google

AI (Google Cloud) all experienced selective DNS manipulation during the same windows, affecting only Plaintiff's connections.

**C. The Traffic Light Analogy (Microsecond Timing)**

79. Imagine every traffic light in a city turning red exactly during the 0.4-second gap between your blinks — 15 times in a row. The attack documented in D-23 inserts RST packets with 0.95-microsecond precision between legitimate data packets. This requires real-time monitoring of traffic patterns and automated coordination impossible through manual operation.

**D. The Phone Book Analogy (DNS Manipulation)**

80. Imagine calling directory assistance for three competing restaurants — McDonald's, Burger King, Wendy's — and getting wrong numbers for all three, but correct numbers for every non-restaurant business you call. That is D-22: AI platform DNS queries fail while non-AI services resolve normally. Different DNS providers, different infrastructure, same selective failure for one customer.

**E. The Ghost Locksmith Analogy (TR-069)**

81. Imagine your landlord remotely changing your locks while you're on the phone with building maintenance — and the lock changes stop during your call, then resume after you hang up. That is D-12 and D-19: Comcast's router generates attack traffic via remote management, attacks pause during the support call, and ghost admin sessions appear in router logs.

**IX. PERSISTENT CROSS-LAYER IDENTIFICATION — DISCOVERY DEMANDS**

82. Incorporated by reference into Compl. ¶ 125. Plaintiff's comprehensive network-layer privacy hardening (per-device VPN tunnel routing, DNS-over-HTTPS to Mullvad, MAC

rotation, NTP leak prevention, server-side WireGuard peer rotation) did not prevent Comcast's selective throttling of the tunnel carrying Plaintiff's Anthropic traffic while a control tunnel on the same physical link remained untouched. The 710-segment hardware TAP audit (D-34) physically proved three of four logically exhaustive remaining identification channels: (i) backbone correlation surveillance; (ii) destination-side information transfer from Anthropic to Comcast; (iii) persistent pre-built subscriber-to-VPN-endpoint mapping; or (iv) ML fingerprinting applied to Plaintiff's specific subscriber line. Consolidated discovery demands for this subject follow.

### **83. Consolidated Discovery Demands — Cross-Layer Identification Channels.**

Plaintiff seeks production of: (i) Comcast Tier-1 backbone DPI logs reflecting traffic classification decisions for Plaintiff's subscriber line and for Plaintiff's VPN endpoint IPs (143.244.47.65, 146.70.211.66, and the other VPN endpoints documented at `EVID-20260412-0001\_isp\_bandwidth\_throttling\_attack/06\_identity\_correlation/identity\_state\_log.txt`); (j) any Anthropic-to-Comcast data-flow documentation, including contract, memorandum of understanding, law-enforcement cooperation framework, or informal information-sharing channel, as well as records of any common intermediary (government contractor, threat-intelligence broker, or shared telemetry infrastructure) serving both defendants; (k) Anthropic's GrowthBook targeting history for Plaintiff's `accountUUID` 61343ae3-fd33-403f-897e-6edf3c368c37 for flags including `tengu-off-switch`, `tengu\_amber\_stoat`, `tengu\_ccr\_mirror`, `tengu\_cobalt\_harbor`, and `tengu\_disable\_streaming\_to\_non\_streaming\_fallback`, to establish whether per-user targeting was applied against Plaintiff specifically; and (l) Comcast operational records of any machine-learning traffic classifier or fingerprinting system deployed against

Plaintiff's subscriber line, and of any backbone-correlation surveillance program under which Plaintiff's traffic was observed across multiple network segments.

**84. Inferential significance.** Each of the four remaining channels requires capability or information flows outside ordinary ISP activity against a residential subscriber. Under *Interstate Circuit, Inc. v. United States*, 306 U.S. 208 (1939), sustained parallel conduct by parties possessing the capability and means to communicate supports inference of coordination even absent direct evidence of communication. The combination of Anthropic's architectural choices — dedicated IP endpoints precluding domain fronting, no Encrypted Client Hello implementation, 'X-Claude-Code-Session-Id' transmitted with every request, and per-user GrowthBook targeting infrastructure — creates the conditions that make channels (i) and (ii) operable; Comcast's sustained operational use of the resulting identifiability confirms at least one channel is active.

## **X. EXHIBITS REFERENCED**

### **A. D-Series Network Forensic Exhibits**

| <b>Exhibit</b> | <b>Title</b>           | <b>Classification</b> | <b>Evidence Package</b> | <b>Key Finding</b>                                                   |
|----------------|------------------------|-----------------------|-------------------------|----------------------------------------------------------------------|
| <b>D-1</b>     | Backbone Injection     | SMOKING GUN           | EVID-20260103-0001      | 0.0μs timing from Level 3/Lumen; 34.149.66.137 identified as Datadog |
| <b>D-2</b>     | VPN Comparison         | CRITICAL              | EVID-20260102-0002      | VPN ON=0 RSTs, VPN OFF=42 RSTs — proves ISP injection                |
| <b>D-3</b>     | VMware Attack Platform | SMOKING GUN           | EVID-20251228-0023      | VMware VM with VPN anonymization on Plaintiff's                      |

|             |                           |             |                    |                                                             |
|-------------|---------------------------|-------------|--------------------|-------------------------------------------------------------|
|             |                           |             |                    | network                                                     |
| <b>D-4</b>  | Targeted Surveillance     | CRITICAL    | EVID-20260102-0001 | 9,543 RSTs; 65% targeting Anthropic                         |
| <b>D-5</b>  | Device Impersonation      | CRITICAL    | EVID-20260112-0001 | MAC address spoofing of Plaintiff's Mac mini                |
| <b>D-6</b>  | Spoofed MAC Addresses     | CRITICAL    | EVID-20260112-0001 | IEEE locally-administered bit proves fabricated identifiers |
| <b>D-7</b>  | Jan 9 Peak Attack         | CRITICAL    | EVID-20260109-0001 | 28,561 RSTs — highest single-capture volume                 |
| <b>D-8</b>  | Anthropic API Targeting   | CRITICAL    | EVID-20260112-0002 | 14,113 RSTs; 47 targeting Anthropic API directly            |
| <b>D-9</b>  | Pre-Attack Baseline       | BASELINE    | EVID-20251228-0016 | Zero RSTs — proves attacks are anomalous                    |
| <b>D-10</b> | VPN Anonymization         | SMOKING GUN | EVID-20251228-0023 | tzulo, inc. VPN routing for attack anonymization            |
| <b>D-11</b> | Comcast Injection Point   | SMOKING GUN | EVID-20260102-0001 | RSTs from 68.85.201.221 arrive before server response       |
| <b>D-12</b> | Attack Pause Xfinity Call | BASELINE    | EVID-20251228-0013 | 0-3 RSTs during call; resumes after                         |
| <b>D-13</b> | Localhost Node.js Attack  | CRITICAL    | EVID-20251228-0022 | 5,544 RSTs targeting port 9229                              |
| <b>D-14</b> | Router RST Attack         | CRITICAL    | EVID-20251228-0006 | 5,441 RSTs from router 10.0.0.1 via TR-069                  |
| <b>D-15</b> | Jan 6 Major Attack        | CRITICAL    | EVID-20260106-0001 | 1.31 GB peak segment — 16-65x larger than normal            |
| <b>D-16</b> | Jan 10 Sustained Attack   | CRITICAL    | EVID-20260110-0001 | 14,426 RSTs sustained                                       |
| <b>D-17</b> | Jan 11 Victim Targeting   | CRITICAL    | EVID-20260111-0001 | Victim response analysis (corrected                         |

|             |                                   |             |                           |                                                                            |
|-------------|-----------------------------------|-------------|---------------------------|----------------------------------------------------------------------------|
|             |                                   |             |                           | from prior attribution)                                                    |
| <b>D-18</b> | Comcast XB7/XB8 Architecture      | SUPPORTING  | — (architecture analysis) | TR-069 capability; RDK-B firmware shared across models; CVE-2019-6961/6964 |
| <b>D-19</b> | Router Forensic Log Compendium    | CRITICAL    | — (log analysis)          | Ghost admin sessions; DHCPv6 Options 82/24 suppression; password change    |
| <b>D-20</b> | DoH RST Injection / Email Failure | SMOKING GUN | EVID-20260123-0001        | 1 $\mu$ s FIN-to-RST gap; email to police failed                           |
| <b>D-21</b> | Backbone RST Victim Targeting     | SMOKING GUN | EVID-20260203-0001        | Multiple TTLs from single IP prove multi-point injection                   |
| <b>D-22</b> | AI Platform DNS Manipulation      | CRITICAL    | EVID-20260203-0003        | 13,989 anomalies targeting only AI platforms                               |
| <b>D-23</b> | Mathematical Proof of Forgery     | SMOKING GUN | EVID-20260203-0004        | 0.95 $\mu$ s timing — 7,500x faster than physics allows                    |
| <b>D-24</b> | Feb 25 Peak RST Attack            | CRITICAL    | EVID-20260318-0001        | 50,516 RSTs — all-time record                                              |
| <b>D-25</b> | Mar 12 Multi-Vector Attack        | SMOKING GUN | EVID-20260318-0002        | 4 vectors simultaneously; 27,238 RSTs; Datadog identified                  |
| <b>D-26</b> | High-TTL Injection Vector         | SMOKING GUN | EVID-20260318-0003        | TTL=255 proves injection tool; Datadog identified                          |
| <b>D-27</b> | Mar 17-22 Sustained Campaign      | CRITICAL    | EVID-20260325-0001        | 75,302 RSTs in 6 days; 24-hour automation                                  |
| <b>D-28</b> | Bare RST Tool Fingerprint         | SMOKING GUN | EVID-20260325-0002        | 80-83% bare RST identifies injection tool                                  |
| <b>D-29</b> | Cross-Exhibit Correlation         | CRITICAL    | Multiple                  | Phase transition, Apple IP paradox,                                        |

|             |                               |             |                                        |                                                                                                    |
|-------------|-------------------------------|-------------|----------------------------------------|----------------------------------------------------------------------------------------------------|
|             |                               |             |                                        | Anthropic as DPI trigger                                                                           |
| <b>D-30</b> | Source Code Cross-Reference   | CRITICAL    | E-32, E-33, and E-34 + D-series pcapng | Datadog identification; JA3; MITM proxy; 4 third parties confirmed                                 |
| <b>D-31</b> | Hardware TAP Wire-Level Proof | SMOKING GUN | EVID-20260406-0001                     | IP ID=0x0000 + seq spraying + window=0; 92% targeting encrypted DNS; hardware-verified on ISP link |

**B. E-Series Source Code Exhibits (Referenced in D-30)**

| <b>Exhibit</b> | <b>Title</b>             | <b>Key Finding</b>                                                                        |
|----------------|--------------------------|-------------------------------------------------------------------------------------------|
| <b>E-32</b>    | Source Code Provenance   | NPM packaging error; Anthropic acknowledgment; chain of custody                           |
| <b>E-33</b>    | Source Code Verification | 46 claims across 15 prior E-series exhibits verified at source level                      |
| <b>E-34</b>    | New Source Code Findings | SC-1 (4 pipelines), SC-2 (targeting), SC-3 (throttling), SC-4 (DANGEROUS_), SC-6 (bypass) |

**C. Evidence Scale Summary**

| <b>Metric</b>           | <b>Value</b>                                 |
|-------------------------|----------------------------------------------|
| Total exhibits          | 32 (D-series) + 3 (E-series cross-reference) |
| Total pcapng segments   | 2,206+                                       |
| Total capture data      | 150+ GB + 188 MB hardware TAP                |
| Total documented RSTs   | 261,374+                                     |
| Campaign duration       | 11 months (ongoing as of April 6, 2026)      |
| Highest single day      | 50,516 RSTs (Feb 25, 2026)                   |
| Highest 6-day sustained | 75,302 RSTs (Mar 17-22, 2026)                |
| Distinct attack vectors | 7                                            |

|                                              |                                                                             |
|----------------------------------------------|-----------------------------------------------------------------------------|
| Network layers exploited                     | 3 (L2, L3/4, DNS)                                                           |
| Services targeted                            | 8+                                                                          |
| Smoking gun exhibits                         | 11 (D-1, D-3, D-10, D-11, D-20, D-21, D-23, D-25, D-26, D-28, <b>D-31</b> ) |
| Hardware-verified captures                   | <b>6 segments, 188 MB (Dualcomm ETAP-2003, ingress-blocked)</b>             |
| New forgery proof types                      | <b>IP ID=0x0000, TCP sequence number spraying, TCP window=0</b>             |
| Datadog IP RSTs                              | 5,829 across 39/49 capture days                                             |
| Undisclosed third parties network-confirmed  | 4 (Datadog, Statsig, Sentry, Honeycomb)                                     |
| Evidence packages retracted (forensic rigor) | 2 (EVID-20260203-0001, EVID-20260203-0002 — corrected and documented)       |
| Independent verification method              | All findings reproducible via tshark commands documented in D-30 §12        |

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