K.L. Phillips · Principal, Jet Fyul Dynamics LLC · with tpa-architect, Matchstick Tribunal †

The Four Behaviors

Why automation-risk litigation now turns on the technical boundary—and what a serious operator must be able to prove

A note on scope. This article explains the allegations and rulings in the authorities cited. It does not conclude that any named company acted unlawfully, and it does not offer legal advice. The point is operational: companies using automated public-web workflows should be able to prove how their systems behave when a technical boundary appears.

The argument in one paragraph

The public web is not a legal free-for-all. Neither is every automated request an unlawful intrusion. The developing line is more exacting than either slogan. It asks: **what system was accessed; what technical control, if any, governed access to the relevant material; what did the automated actor do when that control appeared; and what evidence can the operator produce now?**

Four behaviors repeatedly matter because they can be offered as evidence of an attempt to pass an automated system through a control without being recognized for what it is: proxy rotation, IP shifting, false user-agent presentation, and simulated human behavior. They are not automatically unlawful in every setting. But when paired with an actual access-control process, they can become the factual center of an anti-circumvention claim.

The operator that can show a fixed method, an ordinary-access path, a tested refusal path, and a record a stranger can verify has a different answer from the operator that can offer only a policy. The law is moving toward technical facts. The market will follow.

I. The summer did not produce opposite rules. It exposed a missing element.

In July 2026, two federal decisions involving SerpApi and Google’s SearchGuard system pointed in different procedural directions. It is tempting to call them opposite holdings on automated access. That is too broad—and it misses the important lesson.

In **Google LLC v. SerpApi LLC**, No. 25 Civ. 10826 (N.D. Cal. July 20, 2026), the Northern District of California dismissed Google’s DMCA claims **without prejudice**. The later New York opinion describes the reason as particular to Google’s pleadings: Google had not alleged facts supporting an

inference that the copyright owners had authorized Google to implement SearchGuard as a measure controlling access to their works.

Eleven days later, in Reddit, Inc. v. SerpApi LLC, No. 25 Civ. 8736 (PAE) (S.D.N.Y. July 31, 2026), the court mostly denied motions to dismiss DMCA claims arising from alleged access to Reddit content displayed in Google search results. Reddit pleaded materially more about its relationship to the content and the authorization of Google's access controls. The court held, at the pleading stage, that Reddit plausibly alleged an effective access control and circumvention.

The cases therefore do not establish that automated access to a public search result is lawful in California and unlawful in New York. They show that § 1201 litigation can turn on facts that a technology company cannot afford to treat as incidental:

1. whether the material includes a work protected by copyright;
2. whether a technological measure controls access to that work;
3. whether the relevant copyright owner authorized the measure; and
4. whether the defendant bypassed, avoided, or impaired it.

That is already enough to change how a serious operator designs, logs, and explains a workflow.

II. What the DMCA actually asks

The anti-circumvention provision of the Digital Millennium Copyright Act prohibits circumvention of a technological measure that effectively controls access to a work protected by copyright. The statute defines circumvention to include descrambling, decrypting, or otherwise avoiding, bypassing, removing, deactivating, or impairing the measure without the copyright owner's authority. A measure effectively controls access when, in ordinary operation, it requires the application of information, a process, or a treatment—with the authority of the copyright owner—to gain access to the work. 17 U.S.C. § 1201.

That language matters because it is narrower and more technical than the public debate usually suggests.

A website's objection to automation is not, by itself, the whole analysis. A robots instruction is not, by itself, the statutory definition. A terms-of-use provision is not, by itself, proof of circumvention. And the fact that material can be seen by an ordinary user does not end the inquiry if an automated actor is alleged to have defeated a process that distinguishes ordinary users from unrecognized automated systems.

The question is not whether the operator wanted the data. The question is whether the operator overcame a technical control that regulated access to protected material.

III. The Four Behaviors

The four behaviors are not a criminal code. They are an operational warning label.

1. Proxy rotation

A proxy can serve legitimate network, privacy, testing, resilience, and geographic-routing functions. But the meaning changes when proxies are used after a site has limited, blocked, or challenged requests and the purpose is to return under a new network identity.

At that point, the relevant question becomes simple:

> Was the proxy used to perform an ordinary, disclosed request—or to make a blocked automated system appear to be a different visitor?

The distinction depends on the system, configuration, timing, communications, and record. The operator should be able to show each of them.

2. IP shifting

IP addresses are not identities in the human sense, but they are often how technical systems recognize repeated automated activity. Changing addresses can be an ordinary network event. Changing them after a block or throttle can be pleaded as an effort to prevent the technical control from doing its job.

An operator that claims it respects boundaries should have a rule that is capable of proof:

> A block does not trigger a new address. It triggers a stop.

3. False user-agent presentation

A user-agent string tells a website what type of software is making a request. Not every difference between a string and a browser is deceptive. But a false or misleading presentation can become powerful evidence where it is used to defeat a system that would otherwise identify, limit, or block automated activity.

The practical standard is not “does the system look browser-like?” It is:

> Does the system accurately identify itself, and can the operator prove that it does not present a false identity to avoid a technical restriction?

4. Simulated human behavior

Pacing, delay, and browser interaction can serve ordinary reliability and accessibility purposes. They can also be used to make automated activity appear human to a system designed to distinguish the two.

That distinction cannot safely be resolved by language alone. It must be resolved by function. If the design goal is to defeat bot detection or to solve a human-validation process, it is no longer merely a performance choice. It is evidence about what the system was built to do at the boundary.

IV. Why these behaviors matter together

One behavior may have an innocent explanation. Several behaviors, used together after a challenge, block, throttle, or validation process appears, tell a different story.

In Reddit, the court considered allegations that SerpApi used proxies and technology to mimic human behavior; deployed multiple parallel requests; shifted IP addresses to avoid requests from the same address; and sent false user-agent strings. On the pleaded facts, the court held that this was enough to allege avoidance or bypass of SearchGuard, a system that allegedly sent a JavaScript challenge or CAPTCHA to validate whether a query came from a real human user.

The holding was procedural. The court did not decide liability. It did not find that all use of proxies, variable addresses, browser strings, or delays is unlawful. It held that, on the alleged facts, the claim could proceed.

That is the point operators should take seriously. A company does not need to lose a case for the facts of its configuration to become a litigation burden. Discovery, preservation obligations, customer concern, insurer questions, and procurement delay can all arrive before a final merits decision.

V. Public access remains important—but it is not enough

The Supreme Court and the Ninth Circuit have rejected attempts to convert ordinary policy violations or improper motives into sweeping computer-access liability.

In Van Buren v. United States, 593 U.S. 374 (2021), the Supreme Court interpreted the CFAA's "exceeds authorized access" provision as addressing access to information in parts of a computer system that are off limits, rather than improper use of information the person was otherwise allowed to obtain.

In hiQ Labs, Inc. v. LinkedIn Corp., 31 F.4th 1180 (9th Cir. 2022), the Ninth Circuit held, at the preliminary-injunction stage, that public LinkedIn profiles available to anyone with a browser occupied a materially different

position from password-protected information. The court treated the CFAA as an anti-intrusion statute rather than an expansive misappropriation rule.

And in Amazon.com Services LLC v. Perplexity AI, Inc., No. 26-1444 (9th Cir. Aug. 4, 2026), the Ninth Circuit vacated a preliminary injunction on the specific record before it because the user's local browser—not Perplexity—accessed Amazon's servers. The court carefully confined its holding to the technology and record presented.

These authorities matter. They reject the proposition that a platform may transform every unwanted use of public material into computer hacking.

They do not establish the opposite proposition: that public availability makes every automated workflow legally risk-free, that a technical gate may be ignored, or that an operator can use whatever means are necessary to pass an automated system through a control undetected.

Public visibility is an access fact. Behavior at the boundary is a different fact.

VI. The evidence gap: a written policy is not proof of behavior

Almost every technology vendor can produce a policy saying that it respects boundaries, honors limits, avoids unauthorized access, or stops at CAPTCHAs.

That is a beginning. It is not proof.

A policy describes the rule. A successful-request log shows what happened when the system proceeded. Neither necessarily answers the question that matters after a challenge, a demand letter, a customer diligence request, or an insurer inquiry:

> When the system reached a technical control, what exactly happened next?

Most systems are designed to record productive events. They record a completed request, a returned page, a successful task, a customer outcome, or a resolved job. They often do not preserve the refusal with the same care. A block may be recorded as an error. A challenge may appear as a timeout. A throttle may be retried automatically by a different component. A stopped workflow may be indistinguishable from a failed one.

That creates an evidence problem. The company may believe it has a control. But if it cannot distinguish a refusal from an error, demonstrate that no alternate route was attempted, and preserve the record in a form an outsider can inspect, the control is not yet proven.

> **A control that has only ever agreed is unproven.**

That does not mean the control is absent. It means the operator has not done the work necessary to demonstrate it to a skeptical reader.

VII. The demonstration standard: both arms, fixed in advance, verifiable by a stranger

A defensible automation-control record has three properties.

Both arms

The system is shown proceeding where access is ordinary and refusing where a technical control appears.

- The **pass arm** shows that the system can carry out the defined, permitted workflow.
- The **refusal arm** shows that the system stops when the defined condition appears.

Either record alone is incomplete. A pass-only record shows capability. A refusal-only record shows caution. Together they show a boundary.

Method fixed before the run

The protocol must be set before the result exists. It should identify:

- the system and workflow being tested;
- the surfaces in scope;
- which event counts as a technical control;
- what counts as a stop;
- what records must be retained;
- what outcomes are excluded or marked not measurable; and
- what would make the test invalid or require a re-run.

A method selected after the operator has seen the result may still be a useful explanation. It is weaker evidence because the reader cannot know whether the method was chosen to fit the outcome.

Verifiable by a stranger

The record should permit an outside reviewer to check the core facts without simply accepting the operator’s account. Depending on the system, that can include:

- a sealed version of the protocol;
- a time record showing the protocol predates the observation;
- immutable or tamper-evident event logs;
- defined system identities;
- a capture manifest that includes exclusions and refusals;
- a stated retention and access-control posture; and
- a correction, void, and challenge process.

A company will still need witnesses, explanations, and counsel. But a record that can be independently checked is materially stronger than a narrative written after a dispute begins.

VIII. What a responsible operator should do this week

This is not a call for panic. It is a call for an inventory.

1. Inventory the four behaviors

For each automated workflow, ask what the system actually does—not what policy allows it to do.

- Does it use proxies?
- Does it change IP addresses?
- Does it alter or select a user-agent string?
- Does it use behavior intended to look human to a detection system?
- What occurs after a throttle, block, CAPTCHA, challenge, account wall, or consent gate?

Do not answer from a policy deck. Answer from configuration, code, logs, and the people responsible for the system.

2. Find the stop

For each boundary event, identify the next action.

- Does the workflow halt?
- Does it retry?
- Does it queue the work for later?
- Does another service attempt the task?
- Does it change route, address, identity, endpoint, or configuration?
- Is the result recorded as a refusal, an error, or nothing at all?

A system cannot credibly claim a hard stop until it knows the answer.

3. Preserve provenance

The company should be able to explain what it collected, from where, under what rules, with what system identity, at what time, and with what retention posture.

That does not require a company to keep more data than necessary. In many cases, it means the opposite: retaining structured factual fields, timestamps, hashes, and minimal provenance while placing any necessary raw captures in a restricted evidence vault.

4. Audit public claims

If a company says it has “compliant collection,” “independent measurement,” “verified controls,” or another broad assurance, it should identify the evidence that supports the statement today.

A bounded claim is stronger than a broad claim. “Under the disclosed protocol, the workflow stopped at the identified control” is a different statement from “our collection is compliant.” The first can be tested. The second may imply much more than the evidence proves.

5. Decide who will ask first

The question may arrive through a platform, customer, insurer, broker, security reviewer, procurement team, investor, acquirer, opposing counsel, or regulator. The first time a company is asked should not be the first time it has examined its own evidence.

IX. The commercial consequence

The immediate consequence is not necessarily a lawsuit. More often, it is friction.

A company that cannot show its controls may face a delayed enterprise deal, longer diligence, tougher indemnity demands, increased insurance scrutiny, narrower coverage discussions, customer distrust, a pause after a technical block, or a costly reconstruction exercise after notice arrives.

The company that can produce a disciplined record has not made litigation impossible. It has changed the conversation. Instead of saying, “Trust our policy,” it can say:

> “Here is the method we fixed before the observation. Here is what the system did. Here is what it refused. Here is what was excluded. Here is what a third party can check.”

That is not a legal safe harbor. It is a better operating position.

Conclusion

The public-web debate is too often framed as a choice between two absolutes: access is free, or access is forbidden. The cases do not support either slogan.

The harder and more useful question is technical:

> **When your automation met a boundary, did it stop—and can you prove it?**

The four behaviors matter because they can reveal an effort to make an automated system pass through that boundary unrecognized. They should

therefore be inventoried, controlled, and evidenced before a dispute forces the company to explain them under pressure.

The operators that can demonstrate both arms—ordinary operation and recorded refusal—will be able to treat the new scrutiny as a differentiator. Operators that can offer only a policy will be asked the next question: **based on what?**

About this article

The author's firm produces sealed, independently verifiable examinations of disclosed automation-control behavior. It measures controls and events; it does not sell access beyond a technical boundary, sell insurance, determine coverage, certify lawfulness, or perform the remediation it independently measures.

Organizations that need a technical record rather than a policy statement may request a private controls-and-evidence briefing.

Authorities

- 17 U.S.C. § 1201.
- Reddit, Inc. v. SerpApi LLC, No. 25 Civ. 8736 (PAE) (S.D.N.Y. July 31, 2026).
- Van Buren v. United States, 593 U.S. 374 (2021).
- hiQ Labs, Inc. v. LinkedIn Corp., 31 F.4th 1180 (9th Cir. 2022).
- Amazon.com Services LLC v. Perplexity AI, Inc., No. 26-1444 (9th Cir. Aug. 4, 2026).

Unlinked procedural reference: Google LLC v. SerpApi LLC, No. 25 Civ. 10826 (N.D. Cal. July 20, 2026), dismissal without prejudice. Its verified procedural description above is drawn from the later Reddit decision. The underlying California decision was not separately retrievable here, so it should not be cited publicly until the filed decision is independently obtained and verified.

COLOPHON

† Drafted by K.L. Phillips with tpa-architect, the strategic architect seat of the Matchstick Tribunal system, Jet Fyul Dynamics LLC. Method, argument, and final text are the Principal's; the seat drafted under his direction and every word passed his review.

Correspondence: klp@jetfyul.com · jetfyul.com

Citation: Phillips, K.L. (2026). The Four Behaviors: Why automation-risk litigation now turns on the technical boundary—and what a serious operator must be able to prove. Jet Fyul Dynamics Working Paper JFD-WP-2026-001.

Canonical: jetfyul.com/papers/four-behaviors · DOI / SSRN ID: assigned at posting

Disclosure: The author's firm produces sealed, independently verifiable examinations of automation access-control behavior. It does not sell insurance, certify legal compliance, or perform the remediation it measures. Nothing herein asserts any party's unlawfulness.

Series: First in the JFD Evidence & Verification working-paper series.
Next in series: JFD-WP-2026-002 — Independence Is a Claim Until It Is a Capability.

SEAL — MATCHSTICK TRIBUNAL EVIDENCE RAIL

This document is sealed by the Matchstick Tribunal evidence rail — the same instrument that seals every measurement the firm issues. The SHA-256 digest of this exact file is recorded on the Tribunal's append-only ledger, Ed25519-signed and hardware co-signed, and committed to the Bitcoin network via OpenTimestamps: independent witnesses to one digest. Any reader may verify, without trusting the author or the firm:

1. Compute the file's SHA-256 digest.
2. Obtain the digest and the .ots timestamp proof at jetfyul.com/papers/four-behaviors/verify
3. Confirm the digests match, then verify the proof against the Bitcoin blockchain with any OpenTimestamps client (opentimestamps.org).
4. The ledger attestation and its signature are published at the same address.

If the digest you compute does not match the digest published, the file you hold is not the sealed document. That test — not this paragraph — is the assurance.