

AUGUST 27, 2026 | PUBLICATION

## Client Alert: When an AI Agent Visits a Website, Who Is Really Doing the Accessing? The Ninth Circuit Draws an Early Line Under the CFAA

The Ninth Circuit's decision in *Amazon.com Services, LLC v. Perplexity AI, Inc.* offers one of the first federal appellate examinations of how traditional computer-access laws apply to agentic artificial intelligence (AI) and provides important guidance for businesses deploying AI agents and the platforms those agents access.

As agentic AI becomes increasingly capable of browsing websites, shopping, completing forms, and performing other tasks on a user's behalf, an important legal question has emerged: When an AI agent interacts with a third-party website, who is legally doing the accessing—the user, the AI provider, or the AI agent itself?

On August 4, 2026, the U.S. Court of Appeals for the Ninth Circuit confronted that question in *Amazon.com Services, LLC v. Perplexity AI, Inc.*, No. 26-1444. The court vacated a preliminary injunction that had prohibited Perplexity from using its agentic web-browser technology on Amazon.com and remanded the case for further proceedings.

The decision is significant not because it establishes broad immunity for agentic AI – it expressly does not – but because it begins to define how courts may attribute the actions of AI agents under statutes written long before autonomous software existed.

### The Dispute: Perplexity's AI Agent Goes Shopping

Perplexity AI, Inc. (Perplexity) developed Comet, an AI-enabled web browser that runs locally on a user's computer. Its distinguishing feature is an optional AI agent known as the Assistant, which can perform tasks at the user's direction, including navigating Amazon.com to locate and shop for products.

### INDUSTRY SECTOR

Technology

### SERVICE LINE

Data Privacy  
Technology

### RELATED PROFESSIONALS

C. Jade Davis

### MEDIA CONTACT

Wendy M. Byrne  
wbyrne@shumaker.com

The technical architecture became critical to the Ninth Circuit's analysis.

When a user instructs the Assistant to find an item on Amazon, the Assistant takes screenshots of the browser view and sends them from the user's computer to Perplexity's servers. Perplexity's servers then provide instructions regarding how the Assistant should navigate Amazon. The Assistant therefore relies both on the user's direction and instructions generated through Perplexity's infrastructure.

Amazon objected to this activity. Before Comet's release, Amazon informed Perplexity that its AI products were not permitted to access the Amazon Store. The dispute also involved Perplexity's decision not to use a "user-agent string" that would tell Amazon that a user had activated an AI agent and allow Amazon to block the Assistant.

Amazon sued in November 2025 under the federal Computer Fraud and Abuse Act (CFAA), 18 U.S.C. § 1030, and California's Comprehensive Computer Data Access and Fraud Act (CDAFA), Cal. Penal Code § 502. The district court granted Amazon a preliminary injunction, and the Ninth Circuit vacated it.

#### The Critical Question: Who "Accessed" Amazon's Computers?

The CFAA prohibits certain intentional access to protected computers without authorization or in excess of authorized access. For Amazon's claim under § 1030(a)(2), it therefore was not enough to establish that Amazon had told Perplexity to stay off its platform. Amazon first needed to establish that Perplexity itself "accessed" an Amazon computer within the meaning of the statute.

That distinction proved decisive.

The Ninth Circuit relied on the Supreme Court's interpretation of "access" in *Van Buren v. United States*, which describes access in the computing context as entering a computer system or a particular part of that system.

The court then examined what was actually happening technologically.

Perplexity's servers did not directly communicate with Amazon's servers. Instead, the user's browser communicated with Amazon. The Assistant analyzed information displayed locally through that browser, and information could then be transmitted to Perplexity's servers to determine how to perform the task requested by the user.

That architecture mattered.

The Ninth Circuit concluded that, on the record before it, the user, not Perplexity, was accessing Amazon's computers, albeit with assistance from Perplexity's AI agent. Perplexity's receipt of screenshots and transmission of instructions to the Assistant did not, standing alone, establish that Perplexity had itself gained entry into Amazon's servers.

#### An AI Agent Is a "Tool, Not a Person." at Least Under the CFAA

One of the most consequential portions of the decision concerns the court's treatment of the AI agent itself.

Section 1030(a)(2) applies to "[w]hoever" intentionally accesses a protected computer. The Ninth Circuit reasoned that this language contemplates access by a person or legal entity. The Assistant, regardless of its

sophistication, remained “a tool, not a person for statutory purposes.”

That distinction has implications well beyond this particular dispute.

AI agents increasingly act with varying degrees of autonomy. But statutes imposing liability based on what a person “knows,” “intends,” “accesses,” “obtains,” or “causes” may require courts to determine whose conduct and intent should legally be attributed to the agent’s actions.

The Ninth Circuit did not announce a universal rule of AI attribution. Instead, it focused on the particular technical relationship among the user, the locally operating browser and Assistant, Perplexity’s servers, and Amazon’s servers.

That makes system architecture, not merely the marketing label “agentic AI,” potentially critical to future liability analysis.

#### The Court Was Careful Not to Give Agentic AI a Free Pass

The opinion is notably narrow. The Ninth Circuit expressly declined to decide whether a different factual record could establish that Perplexity exercised sufficient control over the Assistant to itself “gain entry” to Amazon’s servers.

It also emphasized that it was not creating a new legal regime governing agentic AI, nor deciding whether Perplexity could face liability for an AI agent’s conduct under tort law or other legal theories. Its holding was confined to the CFAA’s “access” requirement and the technology reflected in the existing record.

That qualification is important.

Future cases could look very different where an AI provider’s own infrastructure directly communicates with a third party’s servers; the provider independently initiates activity; an agent operates without a specific user instruction; the provider controls credentials or authentication; or the agent circumvents technical access restrictions.

In other words, who operates the infrastructure, who initiates the interaction, and where the actual server-to-server communications occur may matter considerably.

#### The Rule of Lenity Also Constrained the CFAA

The court separately relied on the CFAA’s criminal character.

Although the CFAA provides a civil cause of action in certain circumstances, it is primarily a criminal statute. As a result, the Ninth Circuit applied the rule of lenity in addressing statutory ambiguity.

The court was particularly concerned that Amazon’s proposed interpretation could potentially expose ordinary users themselves to criminal conspiracy or aiding-and-abetting theories merely because they used an AI assistant to interact with a website.

The court cautioned against interpreting the CFAA in a manner that could transform otherwise ordinary computer-assisted behavior into federal criminal conduct.

That reasoning continues a broader trend in CFAA jurisprudence: courts have repeatedly resisted efforts to transform a federal anti-hacking statute into a general-purpose mechanism for policing undesirable uses of

computers or information.

### CDAFA Did Not Change the Result

Amazon also proceeded the CDAFA, which uses language different from the CFAA and can encompass knowingly causing unauthorized computer access.

But the Ninth Circuit concluded that the claims rose and fell together under the circumstances presented. Even assuming California's definition of "access" is broader, the relevant provisions still focus on the person accessing or causing access.

The court therefore reached the same conclusion: the user, rather than Perplexity, accessed Amazon through the Assistant.

### Terms of Service May Become Even More Important

Perhaps the most practically important point for website operators appears near the end of the decision. The Ninth Circuit made clear that its CFAA ruling does not prevent Amazon from regulating access to its platform through private terms of service. The court held only that, on the existing record, Amazon was unlikely to regulate this conduct successfully by invoking the CFAA and CDAFA.

That distinction should matter to virtually every organization operating a public-facing website.

As AI agents increasingly interact with online platforms, businesses should consider whether existing Terms of Use adequately address:

- automated and agentic access;
- AI agents acting on behalf of authenticated users;
- scraping and automated data collection;
- use of bots, crawlers, browser extensions, and autonomous software;
- identification requirements for automated agents;
- credential sharing and delegated authentication;
- use of platform content for AI training or inference;
- circumvention of technological restrictions; and
- contractual remedies for prohibited automated activity.

For many businesses, contract, platform governance, authentication controls, and technical access restrictions may prove more effective than attempting to stretch traditional anti-hacking statutes to cover emerging AI behavior.

### What Companies Deploying AI Agents Should Take Away

The decision also offers an important compliance lesson for companies building or deploying agentic AI.

Developers should understand precisely how their agents interact with third-party infrastructure. An agent running locally in a user's browser presents a materially different legal architecture from an agent that causes the provider's own servers to connect directly to another company's systems.

Organizations should therefore map not simply what their AI agent does, but who technically performs each action: where instructions originate, where credentials reside, which system initiates network requests, which servers communicate directly, what information returns to the AI provider, and how much

independent control the provider exercises after the user's initial instruction.

Those details may ultimately determine whether conduct is attributed to the user, the AI provider, or both.

### The Bigger Picture: AI Attribution Is Becoming a Legal Issue

Perhaps the most important aspect of *Amazon.com Services, LLC v. Perplexity AI, Inc.* is the question it exposes rather than the one it answers.

For decades, technology law generally assumed that a human or company used software to perform an identifiable action. Agentic AI complicates that model. An individual may provide a high-level instruction while software determines the intermediate steps necessary to accomplish it.

As the Ninth Circuit recognized, existing precedent provides imperfect analogies for this emerging technology. The court expressly observed that there is "little to no existing caselaw" addressing how responsibility should be ascribed to AI agents.

Courts will increasingly have to decide when an AI agent's actions should be attributed to its user, its developer, its operator or potentially multiple actors.

For now, *Amazon.com Services, LLC v. Perplexity AI, Inc.* establishes an important early principle: the mere fact that an AI provider's technology helps a user interact with another company's website does not necessarily mean that the AI provider itself has "accessed" that company's computers under the CFAA.

But the Ninth Circuit left the door open for a different result where the architecture, level of control, or factual record changes.

For businesses on both sides of agentic interactions, that means the technical details are no longer merely technical. They may determine who bears legal responsibility for what an AI agent does next.