

Home • Artificial Intelligence

by **Evan Schuman**
Contributor

Federal judge rules for Anthropic in Pentagon dispute, nullifies government supply chain risk designation

News

Aug 28, 2026 • 7 mins

The administration engaged in 'unlawful retaliation in violation of the First Amendment,' the ruling said.



Credit: LightField Studios / Shutterstock

The Trump Administration's decision to punish Anthropic for its stance forbidding Claude's use in domestic surveillance and autonomous weapons by identifying it as a supply chain risk to national security was "arbitrary and capricious," a federal judge ruled on Thursday.

US District Court Judge Rita Lin said federal authorities had no legitimate reason to tell companies with government contracts that they couldn't work with Anthropic.

"The undisputed record shows that the challenged actions constituted unlawful retaliation in violation of the First Amendment and that Anthropic was denied the pre-deprivation process required under the Fifth Amendment," Lin said **in her ruling** [<https://www.courtlistener.com/docket/72379655/250/anthropic-pbc-v-us-department-of-war/>], calling the designation "arbitrary and capricious."

CIO Smart Answers [Learn more](#)

Explore related questions

- [Which defense systems are already integrated with Anthropic's AI technology?](#)
- [Why did the Pentagon label Anthropic a national security supply chain risk?](#)
- [Why did the court call the government's supply chain designation arbitrary?](#)
- [What specific military applications did Anthropic refuse to allow for Claude?](#)
- [What cybersecurity capabilities make Claude Mythos valuable for the U.S. government?](#)

Ask a question

 **ASK**

She stressed that **the government action** [<https://www.cio.com/article/4138822/trump-administration-bans-anthropic-escalating-clash-over-military-use-of-ai.html>] seemed punitive, and was not based on legal and national security risks.

The government's words and deeds "confirm that the challenged actions were based on a desire to make a public example out of Anthropic for its 'arrogance' in criticizing the government, not based on any articulable basis to believe that Anthropic would actually sabotage its model," Lin wrote.

She pointed out, "a few days before the challenged actions began, Secretary Hegseth proposed applying the Defense Production Act to Anthropic, which would mean the company was essential to national security rather than a threat to it. Even now, the

government is discussing collaboration with Anthropic on its new model, Mythos, in an array of sensitive contexts. None of that is consistent with a genuine fear that Anthropic is a saboteur [that] would poison its software to harm national security.”

The judge added that the stated government fears made no sense, noting that the usage policy applicable to Pentagon work is a purely contractual limit. “Anthropic is incapable of enforcing it technologically, and does not have direct visibility into how DoW [Department of War] uses its model,” she pointed out.

“Nothing in the Administrative Record describes, even at a high level, what technological means would give rise to the so-called ‘backdoors’ or could otherwise allow Anthropic to ‘disable’ or affect Claude during a DoW operation,” the judge wrote. “Anthropic has submitted un rebutted evidence that it lacks any technological means to access or control deployed models.”

Lawyers, consultants, and analysts who looked at the decision were confident that the case would be appealed, and that it will end up in the US Supreme Court.

Alan Webber [<https://my.idc.com/getdoc.jsp?containerId=PRF004336>], program VP for national security, defense, and intelligence at IDC, said that Lin’s ruling “was that the label [supply chain risk] was retaliation for Anthropic refusing to loosen safety guardrails DoD [Department of Defense, aka the Department of War] wanted lifted, dressed up in national security language. Put another way, a government customer tried to use a supply chain risk designation as leverage in a contract dispute over model behavior and application, and not because of an actual vulnerability.”

Recommended for you



Trump administration bans Anthropic, seemingly embraces OpenAI

By Cynthia Brumfield • 8 mins


[https://www.cio.com/?post_type=post&p=4138822&utm_source=miso&utm_medium=related&utm_campaign=thumbnail_list]



Anthropic holds firm against the Pentagon after supply-chain risk label

By Anirban Ghoshal • 5 mins

[https://www.computerworld.com/?post_type=post&p=4141684&utm_source=miso&utm_medium=related&utm_campaign=thumbnail_list]



US court refuses to stay Pentagon’s

[https://www.computerworld.com/?post_type=post&p=4156534&utm_source=miso&utm_medium=related&utm_campaign=thumbnail_list]

'supply-chain risk'
blacklisting of
Anthropic
By Gyana
Swain

dium=related&utm_campaign=
thumbnail_list]

• 5
mins

Implications for CIOs

Webber said the implications for CIO strategy are concerning.

"If a government CIO is relying on a vendor's contractual guardrails, this case says those commitments can potentially become the trigger for exactly the kind of blacklisting that risk registers are supposed to protect against," Webber said, noting that anyone who paused Claude usage or froze a subcontract because of the DoD mandate has a legal basis to resume the initiatives. "But obviously that doesn't mean they will, or even should, as this will be appealed."

He added that competing AI vendors have been using the government action as a sales tool, and with this ruling, the argument that Anthropic is a designated supply chain risk "just got weaker, which could lead to contract award disputes."

Consultant **Brian Levine** [<https://formergov.com/directory/brianlevine>], executive director of FormerGov, recommended that CIOs do what they should have always done: Evaluate all products based solely on their merits.

"CIOs should focus on using the frontier models that they believe make the most sense for their business, considering factors such as effectiveness, cost, security, safety, and confidentiality," he said. "Anthropic and the other large frontier models each have too much market share to make retaliation for their use realistic, and the administration seems to have already moved on from this particular battle."

Justin Greis [<https://acceligence.com/talent/profiles/justin-greis/>], CEO of consulting firm Acceligen, agreed that this case has profound implications for CIOs and their AI decisions.

What the federal judge did was reject the leap from a commercial and policy disagreement to an expansive supply chain risk designation without a sufficiently grounded technical rationale or process, **Greis** pointed out.

"The court found that Anthropic did not have the ability to access, alter, or shut down models once deployed in the government environment, and that the government ultimately conceded Anthropic's technology was not inherently riskier than other comparable black box AI models," he said.

"I think that distinction matters enormously for CIOs and CISOs," he stressed. "As AI becomes part of the operating fabric of an enterprise, 'We don't trust the vendor' cannot become a substitute for a defined risk model. Organizations need to be able to

articulate what the actual technical risk is, how it manifests, what controls exist, and whether the response is proportional to that risk.”

“That becomes particularly important with AI,” he added, “because people can easily conflate disagreements over model behavior, usage policies, ethics, contractual restrictions, and cybersecurity into one amorphous category called ‘AI risk.’”

Original government edict still problematic

Mark Rasch [<https://www.linkedin.com/in/raschcyber/>], a former federal prosecutor who is now general counsel at Unit221B, a threat intel and security consulting company, said he was surprised by how quickly government attorneys surrendered on this case. “One of the things that struck me is that the government appears to have abandoned any rationale it might have had for its decision about Anthropic,” he said. The government “came back with all these reasons, but then they abandoned them all when they had to prove them.”

But, he said, the government instruction to all government contractors to also shun Anthropic was problematic.

“It’s one thing for the government to say ‘We’re not going to do business with you.’ It’s quite another thing to say ‘Nobody we do business with can do business with you either,’” Rasch said. “This says that if you are disfavored by the administration, they’re not just going to blacklist you and say they won’t do business with you. They’re going to say that nobody can do business with you.”

Supreme Court arguments will likely be very different

Rasch predicted that the legal arguments in the Supreme Court will be quite different, and will potentially sidestep the lack of evidence.

“In the Supreme Court, [the government’s] biggest argument will not be that ‘We are right that it is a supply chain risk,’ but that, ‘Whether we’re right or wrong is irrelevant. We get to make that [supply chain risk designation] decision, not the court.’”

That would mean that the Supreme Court Justices could avoid exploring whether the government made the right decision, and instead focus on whether the government has the unlimited right to decide who is a national security risk.

*This article originally appeared on **Computerworld** [<https://www.computerworld.com/article/4215393/federal-judge-rules-for-anthropic-in-pentagon-dispute-nullifies-government-supply-chain-risk-designation.html>].*

[Artificial Intelligence\[https://www.cio.com/artificial-intelligence/\]](https://www.cio.com/artificial-intelligence/)

[Government\[https://www.cio.com/government/\]](https://www.cio.com/government/)

[Industry\[https://www.cio.com/industry/\]](https://www.cio.com/industry/)

[Legal\[https://www.cio.com/legal/\]](https://www.cio.com/legal/)

[Business Operations\[https://www.cio.com/business-operations/\]](https://www.cio.com/business-operations/)

[Vendors and Providers\[https://www.cio.com/vendors-and-providers/\]](https://www.cio.com/vendors-and-providers/)

NEWSLETTER

From our editors, straight to your inbox

News, analysis, and insights for IT leaders navigating the risks and rewards of AI. A special series from the editors of CIO, Computerworld, CSO, InfoWorld and Network World



By submitting your information, you agree to our [PRIVACY POLICY \[https://foundryco.com/privacy-policy/\]](https://foundryco.com/privacy-policy/).

SUBSCRIBE

Sponsored Links

See Why Five9's AI-Powered CX Platform Drives Real Results

From pilots to performance: AI-driven growth with EY Consulting.

IDC Examines Rocket® EVA™ and Agentic AI for IT Operations

Is legacy tech holding back enterprise AI value? Tune in for insights.

The fastest PCs with multi-day battery life for your business.

About

[About Us](#)

[Advertise](#)

[Contact Us](#)

[Editorial Ethics Policy](#)

[Foundry Careers](#)

[Reprints](#)

[Newsletters](#)

[Add CIO as a Preferred Source in Google Search](#)

More

[Awards](#)

[Blogs](#)

[BrandHubs](#)

[BrandPosts](#)

[CXO Outlook](#)

[DEMO](#)

[Buyer's Guides](#)

[Events](#)

[Podcasts](#)

[Videos](#)

Policies

[Terms of Service](#)

[Privacy Policy](#)

[Cookie Policy](#)

[Copyright Notice](#)

[Member Preferences](#)

[About AdChoices](#)

[Your California Privacy Rights](#)

Our Network

[Computerworld](#)

[CSO](#)

[InfoWorld](#)

[Network World](#)

FOUNDRY



