

Vital Briefing

Anthropic's Dispute With the Department of War

By [George M. Gullo, J.D.](#)

In late February 2026, the United States government took an extraordinary step: in a social media post President Trump ordered all federal agencies to cease using technology from Anthropic, the American AI company behind the Claude model. Department of War* (DoW) Secretary Pete Hegseth followed that by formally designating Anthropic a “supply chain risk to national security.” This designation—the first ever applied to a domestic company—stems from a contractual dispute over whether Anthropic must allow the military to use its AI models for all lawful purposes without restriction, including in fully autonomous weapons and mass domestic surveillance. Anthropic refused and on March 9, Anthropic sued the Trump administration in two federal courts.

Background and Contractual Dispute

In July 2025, Anthropic entered into a contract with the Chief Digital and Artificial Intelligence Office, an agency within the Department of Defense, valued at up to \$200 million over two years—structured as a ceiling contract with usage-based payments rather than a lump-sum purchase, meaning only a portion of the ceiling value had been paid out before the dispute arose. The agreement made

Anthropic the first AI company to have its model approved for deployment on classified government networks. As part of the deal, the DoW agreed to abide by Anthropic's Acceptable Use Policy (AUP), which contained two notable restrictions:

- Claude could not be used to create or operate fully autonomous lethal weapons systems without human oversight.
- Claude could not be used for mass domestic surveillance of American citizens.

These restrictions reflected Anthropic's position that current AI models are not yet dependable enough for unsupervised lethal decision-making, and that mass surveillance represents a fundamental threat to democratic governance. Anthropic has stated publicly that these two restrictions “have not affected a single government mission to date.”

“Any Lawful Use” Demand

In January 2026, Secretary Hegseth issued a memorandum setting a new “AI-first” agenda for the DoW that directed all AI contracts to incorporate standard “any lawful use” language within 180 days. Applied to the Anthropic contract, this

would have required Anthropic to remove its two AUP restrictions, allowing use of Claude for autonomous weapons and domestic surveillance programs. After weeks of negotiations, the DoW set a hard deadline of 5:01 p.m. on February 27, 2026, for Anthropic to comply. Anthropic declined.

Anthropic's statement defended its position: “We cannot in good conscience accede to their request.” The company noted the inherent tension in the government's posture—simultaneously threatening to invoke the Defense Production Act to compel use of Claude on grounds it is essential to national security, while also designating it a supply chain risk. Anthropic CEO Dario Amodei acknowledged the company had never raised objections to specific military operations. The company's position was that the contractual restrictions against fully autonomous weapons and mass surveillance were the appropriate place to draw the line—not ad hoc intervention in operational decisions.

Presidential Order and Designation

On February 27, 2026, President Trump posted on Truth Social directing all federal agencies to immediately cease

* The Executive Branch has announced its intention to rename the Department of Defense the “Department of War.” While that rename has not been effectuated by statute as of this writing, both federal court captions in this litigation—*Anthropic PBC v. U.S. Department of War*, No. 3:26-cv-01996 (N.D. Cal.), and *Anthropic PBC v. U.S. Department of War*, No. 26-1049 (D.C. Cir.)—use “Department of War” for the defendant's name. Amicus filings from the ACLU and 149 former judges similarly caption the cases using that name. This briefing follows the court captions accordingly. Where formal statutory text or agency designations require reference to the Department of Defense, that name is used.

using Anthropic's technology, with a six-month phase-out period for most agencies. Secretary Hegseth simultaneously declared Anthropic a "Supply-Chain Risk to National Security" on X (formerly Twitter). Neither directive initially cited specific legal authority.

The formal designation letter, dated March 3, 2026, confirmed the supply chain risk designation as applying to all Anthropic affiliates and all Anthropic products and services used in connection with DoW contracts. The GSA followed suit and removed Anthropic from USAi.gov, the government's centralized platform for testing AI models. In a television interview, DoW Chief Technology Officer (CTO) Emil Michael subsequently stated that Claude would "pollute" the defense supply chain due to its "different policy preference" embedded in its design.

Despite the designation, the DoW reportedly continued using Anthropic's Claude on classified military networks—through Palantir's Maven Smart System—during the U.S. military campaign in Iran, hours after announcing the ban. The Department acknowledged it would take time to transition and could not "just rip out" Anthropic's technology overnight and planned a six-month transition. Palantir CEO Alex Karp confirmed his company continued to use Claude.

Legal and Regulatory Framework

The DoW's designation purports to rest on two distinct statutory frameworks that operate under different rules and are subject to different judicial review processes, which is why Anthropic filed lawsuits in two separate courts.

10 U.S.C. § 3252 / DFARS Subpart 239.73. This provision permits the Secretary of Defense to exclude companies from "covered procurement actions" involving "covered systems," which are defined as National

Security Systems (NSS). The statute was designed to address foreign adversary supply chain threats and requires a specific finding that "an adversary may sabotage, maliciously introduce unwanted function, or otherwise subvert" a covered system. The government is likely to argue that the Secretary retains broad discretion in making that finding and that the statute does not require identification of a foreign state actor. Critics contend, however, that Anthropic is an American company and the underlying dispute is contractual disagreement, not adversarial sabotage. Critically, the statute's reach is limited to NSS contracts. It does not extend to routine administrative systems, and explicitly excludes payroll, finance, logistics, and personnel management applications.

Federal Acquisition Supply Chain Security Act of 2018 (FASCSA), 41 U.S.C. §§ 1321–1328; Exclusion Authority, 41 U.S.C. § 4713. FASCSA provides broader authority, potentially excluding a designated source from all federal contracts—not just DoW NSS contracts—and directing contractors not to use the designated source in their subcontracts (implemented through FAR 52.204–30). The FASCSA process provisions are codified at 41 U.S.C. §§ 1321–1328; the exclusion order authority itself is codified at 41 U.S.C. § 4713. A FASCSA order would require the Federal Acquisition Security Council to complete a supply chain risk assessment, consider less intrusive alternatives, and provide notice to the affected source before issuing a recommendation. Based on Anthropic's court filings, those procedural prerequisites do not appear to have been satisfied.

FAR implementing clauses. The Federal Acquisition Regulation contains two FASCSA-implementing clauses with contractor implications: FAR 52.204–29 (Prohibitions and Representations) governs supply chain security representations and prohibitions;

FAR 52.204–30 (Reporting and Monitoring) imposes prompt monitoring and reporting obligations if a formal FASCSA order issues against Anthropic. Also, FAR Subpart 9.4 (Debarment and Suspension), if invoked, could bar Anthropic from receiving new contract and subcontract awards.

DFARS implementing clauses. Two Defense Federal Acquisition Regulation clauses implement the supply chain risk framework at the contract level: DFARS 252.239–7017 (Notice of Supply Chain Risk) notifies offerors that the government may exercise supply chain risk authorities under 10 U.S.C. § 3252, including consideration of intelligence information; and DFARS 252.239–7018 (Supply Chain Risk) requires contractors to "mitigate supply chain risk in the provision of supplies and services to the Government," but applies only to covered defense contracts involving NSS and does not prohibit commercial business relationships with Anthropic outside of specific defense work.

Anthropic's Grounds for Legal Challenges

On March 9, 2026, Anthropic filed lawsuits in two forums because the two statutory bases for the designation—10 U.S.C. § 3252 and FASCSA—are reviewed under different frameworks.

In the Northern District of California (No. 3:26-cv-01996), Anthropic sought a temporary restraining order (TRO) challenging the § 3252 designation under the Administrative Procedure Act (APA) and on constitutional grounds. That case names 18 federal agencies and 17 officials as defendants across five counts:

- violation of 10 U.S.C. § 3252 and the APA;
- First Amendment retaliation;
- ultra vires presidential action;
- Fifth Amendment due process; and
- unlawful agency sanctions.

In the Court of Appeals for the District of Columbia Circuit (No. 26-1049), Anthropic challenged the FASCSA-based aspects of the designation and sought an emergency stay pending review. The D.C. Circuit is the statutorily prescribed forum for FASCSA challenges. FASCSA's judicial review provision grants the D.C. Circuit original jurisdiction to hear petitions from companies designated under 41 U.S.C. § 4713, rather than routing the challenge through a district court as would be typical for APA claims. This direct appellate jurisdiction reflects Congress's recognition that FASCSA exclusions are of national significance and warrant expedited review at the circuit level.

Anthropic's complaints identify several grounds for challenge, which legal analysts have assessed as follows:

Statutory misfit. The § 3252 designation requires a finding that “an adversary” may subvert covered systems. Anthropic contends it is an American company, not an adversary, and that the dispute is a contract negotiation, not a security threat. Applying a statute designed for foreign adversary infiltration—such as Huawei, ZTE—to a domestic company over a policy disagreement presents a textual question that courts will be asked to scrutinize closely.

Major questions doctrine. The major questions doctrine, as articulated by the Supreme Court in *West Virginia v. EPA* (597 U.S. 697) and reaffirmed in *Biden v. Nebraska* (600 U.S. 477), holds that when an agency claims authority to resolve questions of vast economic and political significance, courts will not defer to that interpretation unless Congress has spoken clearly to authorize it. The doctrine is a check on agencies that locate sweeping powers in ambiguous statutory text.

Anthropic argues that the government is using FASCSA—a statute enacted to protect federal procurement networks from foreign adversary hardware and software infiltration—to accomplish something outside the statute's original scope: conditioning a domestic AI company's participation in federal contracting on the removal of safety constraints from its commercial product. Anthropic contends that it is a question of vast economic and political significance. The AI industry's commercial relationships with the federal government add up to billions of dollars annually. A ruling that the executive branch can condition those contracts on the removal of any “policy preference” embedded in a product's design would give agencies broad leverage over the design choices of every technology company that sells to the government.

The major questions doctrine also intersects with the First Amendment argument. Some scholars argue that an AI model's safety constraints are a form of expressive design—that a company's decision about what its product will and will not do reflects protected speech or association. If that argument gains traction, the major questions doctrine provides additional support: Congress would need to speak with unmistakable clarity before authorizing the government to override such choices, and FASCSA plainly does not do so.

Procedural defects. FASCSA requires the Federal Acquisition Security Council to complete a risk assessment, consider less intrusive alternatives, and provide prior notice to the affected source before issuing a recommendation. These procedural prerequisites appear to have been bypassed entirely. Secretary Hegseth announced the designation on social media before any formal administrative record was compiled,

which Anthropic argues constitutes arbitrary-and-capricious agency action under the APA.

First Amendment concerns. Some scholars have raised the question of whether requiring a company to remove safety restrictions as a condition of contracting implicates expressive or associational rights. DoW CTO Emil Michael's statement that Claude would “pollute” the supply chain due to its “different policy preference” is notable in this regard: it is an on-record statement by a senior official that the government's objection concerns the content of the company's design choices rather than an operational security threat in the conventional sense.

Internal contradiction. Anthropic argues that the DoW cannot simultaneously maintain that Claude poses an acute supply chain threat requiring emergency exclusion while also deeming it safe to continue using Claude for six months in active combat operations. The government's own conduct—continuing to deploy Claude through Palantir's Maven Smart System during the Iran campaign—is expected to bear directly on the irreparable harm and public interest factors in preliminary injunction analysis.

Government's Rebuttal Arguments

National security deference and statutory authorization. The government's strongest argument is the breadth of executive deference in national security procurement. Under *Dept. of the Navy v. Egan* (484 U.S. 518) and its progeny, courts have consistently declined to second-guess executive branch determinations about what poses a threat to national security systems. The designation here was made through statutory mechanisms—FASCSA and § 3252—that Congress specifically created

for supply chain risk management. Courts are far more reluctant to intervene when the executive acts within an explicit statutory grant than when it is stretching ambiguous text.

Misapplication of major questions

doctrine. On the major questions doctrine, the government has a credible argument that the doctrine is being misapplied. *West Virginia v. EPA and Biden v. Nebraska* targeted agencies claiming sweeping regulatory powers hidden in minor statutory provisions—the “elephant in a mousehole” statutory interpretation problem (see *Whitman v. American Trucking Assns., Inc.*, 531 U.S. 457). FASCSA is a dedicated statutory scheme with an inter-agency committee, defined procedures, and explicit exclusion authority. The government can plausibly argue that using a clearly authorized tool to exclude a clearly covered article is not the kind of novel government overreach to which the doctrine applies, regardless of the large size of the affected market.

Structural limits of First

Amendment claim. Companies have no constitutional right to contract with the federal government, and the government retains broad authority to set conditions on who it buys from and what those products must do. The unconstitutional conditions doctrine requires that the government not demand the relinquishment of a constitutional right as the price of a benefit. The government here can argue that requiring operationally reliable, unfiltered AI outputs from a combat system is a legitimate product specification, not viewpoint-based censorship. Framed that way, it looks more like a standard procurement requirement than a First Amendment violation.

Emergency authority justifies procedural

departure. On the procedural defects, the government can invoke emergency national security authority to justify bypassing standard FASCSA notice and risk assessment requirements. Courts have historically permitted post-hoc administrative record compilation in genuine national security emergencies, and the government will argue this falls within that tradition. Whether that argument survives scrutiny depends heavily on whether the court accepts the “emergency” framing—which is exactly what Anthropic is contesting.

Amicus Briefs

Amicus support for Anthropic has been filed across a broad and institutionally diverse range of organizations and individuals, spanning the technology industry, civil liberties organizations, conservative legal institutions, the national security community, and the federal judiciary. The breadth of the briefs reflects a range of interests that regard the outcome of the litigation as significant beyond the immediate parties.

On the same day Anthropic filed its lawsuits, a group of 37 individuals identifying themselves as “engineers, researchers, scientists, and other professionals” in the AI industry filed an amicus curiae brief in the Northern District of California supporting Anthropic’s TRO motion. All signed in their personal capacities, expressly not as representatives of their employers. Also among the amici were individuals affiliated with OpenAI and Google DeepMind. Additional amicus briefs were filed in the district court preliminary injunction proceedings by Microsoft; the ACLU (jointly with the Center for Democracy & Technology); the Cato Institute; the

Chamber of Progress; FIRE and the First Amendment Lawyers Association; former senior national security officials; and a coalition of Catholic moral theologians and ethicists.

The AI employee brief advances three principal arguments that bear on the merits:

The designation was improper and

arbitrary. The brief states that the government’s designation of Anthropic as a supply chain risk “was an improper and arbitrary use of power that has serious ramifications for our industry.” The signatories argue that the government had an obvious, less punitive alternative—canceling the contract and moving to another provider—rather than deploying a national security blacklist against an American company over a contract disagreement.

The red lines are technically and ethically

legitimate. The brief states directly that “the best currently available AI systems cannot safely or reliably oversee fully autonomous lethal targeting and should not be available for domestic mass surveillance of the American people.” This is a concession by researchers at competing firms that Anthropic’s technical objections are sound.

Developer safety guardrails are the last

line of defense. The brief argues that “in the absence of public law, the contractual and technological requirements that AI developers impose on the use of their systems represent a vital safeguard against their catastrophic misuse.” There are currently no federal statutes governing AI use for autonomous weapons or mass surveillance, making contractual restrictions the only formal protection in place.

The amicus support for Anthropic has proven far broader than the initial filing by the industry employees. On March 17, 2026, a bipartisan coalition of 149 former judges filed an amicus brief supporting Anthropic's emergency stay motion in the D.C. Circuit, the forum in which Anthropic challenges the FASCSA-based aspects of the designation. That brief is particularly significant given the D.C. Circuit's central role in administrative law and its broad jurisdiction over federal agency action. The former judges argue that DoW misinterpreted the governing statute and failed to follow mandatory procedural prerequisites, and that courts have both the authority and the obligation to review such executive action without second-guessing underlying national security judgments. According to the judges, the designation process violated basic principles of due process and administrative regularity, and the absence of the statutory procedural prerequisites (risk assessment, consideration of less intrusive alternatives, and prior notice) raises significant questions about the scope of judicial review.

Preliminary Injunction Hearing and Ruling

A preliminary injunction hearing before Judge Rita F. Lin in the Northern District of California took place on March 24, 2026. The hearing was notable for the bench's skepticism toward the government's position. Judge Lin characterized the designation as "troubling" and described it as "an attempt to cripple Anthropic," observing that the national security concerns identified by the government could have been addressed simply by ceasing to use Claude rather than invoking supply chain risk designation authority. She pressed government counsel on the absence of a meaningful limiting principle, noting that the government's theory appeared broad

enough to permit designation of virtually any vendor that insisted on particular contract terms—"a really low bar" that could apply in any case.

Two days after the hearing, Judge Lin granted Anthropic's motion for a preliminary injunction to temporarily pause the government's supply chain risk designation and related adverse actions against Anthropic. The ruling allows Anthropic to continue operating and conducting business with federal agencies and contractors while the case proceeds on the merits. Judge Lin stayed implementation of the order for one week to permit the government to seek a stay from the Court of Appeals for the Ninth Circuit before the injunction takes effect.

Judge Lin grounded her ruling on both the First Amendment and the Fifth Amendment Due Process Clause, which addressed two of the five counts alleged in Anthropic's complaint. The key language from the order on the First Amendment count is: "Punishing Anthropic for bringing public scrutiny to the government's contracting position is classic illegal First Amendment retaliation." On the statutory misfit argument, Judge Lin wrote: "Nothing in the governing statute supports the Orwellian notion that an American company may be branded a potential adversary and saboteur of the U.S. for expressing disagreement with the government."

The First Amendment retaliation holding is legally significant for several reasons. First, it is the cleanest and most portable of the available doctrinal grounds—it does not require the court to resolve the contested statutory interpretation questions under 10 U.S.C. § 3252 or FASCSA, nor does it require the court to reach the major questions doctrine analysis. Second, a First Amendment retaliation finding at the preliminary injunction stage requires a

showing of a likelihood of success on the merits, meaning Judge Lin concluded that Anthropic is likely to prevail on this claim at trial. Third, First Amendment retaliation doctrine is well-established and appellate courts have developed clear standards, which may make this ground more durable on appeal than the statutory arguments. The government's rebuttal argument—that Anthropic's refusal to accept contract terms constitutes commercial conduct rather than protected speech—was not persuasive to the court at this stage.

The injunction does not require the government to use Anthropic's technology. As Judge Lin observed during the hearing, the Department of War is free to stop using Claude and seek a more permissive AI vendor. What the injunction prohibits is the designation itself and the related actions that flowed from it pending resolution of the litigation.

Conclusion

The Anthropic–Department of War dispute is a consequential AI procurement controversy. It raises fundamental questions about the limits of national security procurement authority, whether executive agencies can use supply chain security statutes to pressure domestic technology companies on matters of product design, and ultimately who controls the ethical constraints on military AI. The designation of an American AI company as a national security supply chain risk—a tool previously reserved for foreign adversaries—is legally unprecedented. Ultimately, the outcome of Anthropic's litigation will likely determine whether the U.S. military can mandate the removal of private safety guardrails or if the courts will narrow this authority to prevent what Anthropic alleges is an "unlawful campaign of retaliation" against its protected corporate speech.