

First Take: Anthropic-Pentagon Rift Signals a Sovereign AI Reckoning

28 February 2026 - ID G00850419 - 5 min read

By: Patrick Connaughton, Yanni Karalis, Lydia Clougherty Jones, Adam Preset, Gary Olliffe
Initiatives: Technologies and Markets; Maximize the Return on AI Investments;
Accelerating Enterprise AI Value Realization at Scale; Strategy, Risks and Opportunities

The Pentagon has designated Anthropic a supply chain risk. It has ordered that no contractor, supplier or partner doing business with the U.S. military conduct any commercial activity with Anthropic. This may prompt a significant ripple effect in the industry. The recent dynamics highlight the necessity for building contractual and operational resilience as AI vendors face increasing geopolitical and regulatory scrutiny.

Anthropic-Pentagon Talks Signal New Era in AI Ethics and Security

Claude is extensively deployed across the U.S. Department of War and other national security agencies for mission-critical applications, such as intelligence analysis, modeling and simulation, operational planning, cyber operations, and more. For example, the U.S. military used Anthropic's Claude AI models to analyze intelligence and process satellite images during the operation to capture former Venezuelan President Nicolas Maduro. ¹

Anthropic and the Pentagon have been in ongoing discussions about the terms of use for Claude in military operations. Anthropic has maintained a firm stance against its AI being used for mass surveillance of Americans or in fully autonomous weapons systems, while the Pentagon would like to remove these restrictions to allow for military operations.

This incident serves as a broader stress test for the evolving balance of global influence. It goes beyond a single agreement, highlighting a clash between corporate responsibility, national security imperatives and an emerging world order shaped by sovereign AI. Technology leaders should prioritize centralized data and AI controls and build flexible systems that can quickly switch providers to stay resilient amid regulatory and geopolitical shifts.

Notably, OpenAI has already agreed to allow use for unclassified activities and is open to intelligence work and battlefield missions under lawful terms. As of mid-January, xAI's Grok continues to expand its integration into both unclassified and classified Pentagon networks. ²

On 26 February, Anthropic officially released a statement including these key take-aways: ³

- Anthropic has never raised objections to particular military operations nor attempted to limit use of their technology in an ad hoc manner.
- However, Anthropic's CEO states that some uses are also simply outside the bounds of what today's technology can safely and reliably do – specifically, mass domestic surveillance and fully autonomous weapons.

On 27 February, the administration ordered federal agencies and contractors that work with the military to immediately cease business with Anthropic. Some government agencies, including the Pentagon, have six months to phase out use of Anthropic's products. ⁴

As a result, Anthropic may lose its government contract and access to commercial partners that do business with the Pentagon. This could have a significant ripple effect, as major providers such as AWS and Palantir, which hold billions in military contracts, may need to reassess their relationships with Anthropic.

It's important to note that this designation is typically reserved for non-U.S.-owned entities whose products or services could enable an adversary to sabotage, subvert or degrade critical systems. Applying it to a U.S. company like Anthropic is unprecedented, particularly since the issue centers on refusal to remove certain guardrails, not a specific security incident.

The Pentagon invoking the U.S. Defense Production Act (DPA) could compel Anthropic to lift current restrictions on the use of its model, regardless of the current contract terms to the contrary, as the DPA allows the U.S. government emergency authority to take control of critical private-sector assets under certain circumstances. These include ensuring the availability of industrial resources and services for national security and defense (e.g., military, energy and emergency preparedness activities).⁵

However, triggering the DPA is not a *fait accompli* as it is subject to the Pentagon defending its extension to the current situation, alongside objection by Anthropic and others, including that the new terms cannot reliably be performed.

While the risk designation and/or invoking the DPA are seemingly inconsistent – the former requiring a security risk and the latter being essential to security – they are fact-based strategic levers, enabling both sides to position applicability/nonapplicability and strengthen or prolong negotiations.

A possible compromise could involve Anthropic providing a defense-specific, isolated and secure version of Claude, allowing lawful use while maintaining transparency and human-in-the-loop oversight. This could set a precedent for how AI labs balance private ethics with government requirements.

Implications for IT Leaders

No single AI vendor is too big to fail. Organizations should ensure they can swap out models without disrupting operations. Commercial clients with federal revenue streams will have to choose between Anthropic and government contracts.

As a best practice, all organizations should decouple applications from specific AI providers – such as Anthropic, Google or OpenAI – using LLM gateways, abstraction layers or platform choices that enable switching models with minimal disruption. This reduces vendor lock-in and mitigates risk as the AI landscape evolves.

Implications for IT Sourcing and Procurement

The recent dynamics with Anthropic highlight the necessity for building contractual and operational resilience as AI vendors face increasing geopolitical and regulatory scrutiny.

For procurement and sourcing leaders, ensure contracts allow for immediate termination without penalty if a risk designation occurs. Include regulatory contingencies in AI service agreements to protect organizational interests.

Start with contracts that include strong termination and contingency clauses, enabling immediate exit without penalty if a vendor receives a government-imposed supply chain risk designation or restriction.

IT sourcing, procurement and vendor management must require data portability, multimodel evaluation and vendor lock-in protections to avoid dependence on a single AI provider. Vendor risk assessments should expand to cover geopolitical exposure, federal revenue dependence and compliance vulnerabilities, ensuring risks are understood and documented well before any disruption. This will also provide resilience against the effects of direct or indirect DPA invocation. Regarding these topics, leaders responsible for IT sourcing and procurement should collaborate with their legal counsel on contract strategy.

By strengthening contracts and improving vendor intelligence, IT sourcing and procurement leaders can keep the organization compliant, avoid operational disruption and maintain both control and strategic flexibility in an increasingly geopolitically volatile AI landscape.

Evidence

¹ [Pentagon Used Anthropic's Claude in Maduro Venezuela Raid](#), The Wall Street Journal.

² [Musk's AI Tool Grok Will be Integrated Into Pentagon Networks, Hegseth Says](#), The Guardian.

³ [Statement from Dario Amodei on our discussions with the Department of War](#), Anthropic.

⁴ [Hegseth Says Pentagon Designating Anthropic as Supply Chain Risk After Trump Bans AI Firm](#), The Hill.

⁵ [What Is the Defense Production Act?](#), Council on Foreign Relations.

Evolving Insights History

Table 1: Evolving Insights History

<i>Published Date</i> ↓	<i>Update</i> ↓
28 February 2026	On 27 February, the Pentagon released a statement saying that, effective immediately, no contractor, supplier or partner that does business with the United States military may conduct any commercial activity with Anthropic.
27 February 2026	Major update: On 26 February, Anthropic officially released a "Statement from Dario Amodei on our discussions with the Department of War." This update accounts for that announcement, specifically Anthropic's position that "we cannot in good conscience accede to their request."
19 February 2026	Initial Publish

Recommended by the Authors

Some documents may not be available as part of your current Gartner subscription.

[AI Solution Report: Anthropic Claude](#)

[First Take: With Opus 4.6, Anthropic Sets Its Sights on Automating Business Workflows](#)

[Top Trends in D&A for 2026: Global Sovereign AI Accelerates](#)

[Sovereign AI Series: Shift Your AI Gravity to Regional Autonomy](#)

[IT SPVM Leaders: 3 Approaches to Address AI Vendor Lock-In Risk](#)

© 2026 Gartner, Inc. and/or its affiliates. All rights reserved. Gartner is a registered trademark of Gartner, Inc. and its affiliates. This publication may not be reproduced or distributed in any form without Gartner's prior written permission. It consists of the opinions of Gartner's Business and Technology Insights Organization, which should not be construed as statements of fact. While the information contained in this publication has been obtained from sources believed to be reliable, Gartner disclaims all warranties as to the accuracy, completeness or adequacy of such information. Although Gartner insights may address legal and financial issues, Gartner does not provide legal or investment advice and its insights should not be construed or used as such. Your access and use of this publication are governed by [Gartner's Usage Policy](#). Gartner prides itself on its reputation for independence and objectivity. Its insights is produced independently by its Business and Technology Insights Organization without input or influence from any third party. For further information, see "[Guiding Principles on Independence and Objectivity](#)." Gartner insights may not be used as input into or for the training or development of generative artificial intelligence, machine learning, algorithms, software, or related technologies.

Table 1: Evolving Insights History

<i>Published Date</i> ↓	<i>Update</i> ↓
28 February 2026	On 27 February, the Pentagon released a statement saying that, effective immediately, no contractor, supplier or partner that does business with the United States military may conduct any commercial activity with Anthropic.
27 February 2026	Major update: On 26 February, Anthropic officially released a "Statement from Dario Amodei on our discussions with the Department of War." This update accounts for that announcement, specifically Anthropic's position that "we cannot in good conscience accede to their request."
19 February 2026	Initial Publish