Monday, March 2, 2026

tomorrowsaffairs.com

TA Tomorrow's Affairs

Analysis of today
Assessment of tomorrow

By: Tomorrow's Affairs Staff

How the Pentagon banned the artificial intelligence it used to wage war

Between 27 February and 1 March, the United States Department of Defense issued a decision classifying a commercial artificial intelligence as a national security risk and formally excluding it from federal use.

At the same time, the same system, through existing operational arrangements, remained active on classified military networks during ongoing combat operations.

This sequence of events was not the result of an administrative error but the consequence of a structural mismatch between political decision-making and the system's operational dependence on technology that was already deeply integrated.

In public reports, the mismatch was simplified as a conflict between the federal government and a technology company.

In reality, it demonstrates how a state can formally suspend a tool without immediately suspending its use, thereby avoiding compromise to its own operational capabilities.

In this context, the simultaneous prohibition and use of the same technology is not an exception but a functional reality of modern warfare.

How Claude entered the US military system

Understanding the events of February requires returning to January 2026. At the beginning of that month, US special forces conducted an operation on Venezuelan territory that resulted in the arrest of President Nicolás Maduro.

A total of 83 people died during the operation, including 47 members of the Venezuelan armed forces.

The Wall Street Journal later reported that the Claude artificial intelligence system, developed by Anthropic, was used in the preparation and execution of the operation through the Palantir software platform, which is employed in the most sensitive military and intelligence structures of the United States.

Anthropic denied having direct communication with Palantir regarding the operation, while the Department of Defense made the opposite claim.

Regardless of these disputed allegations, the fact remains that Claude was integrated into the Department of Defense's operational systems at that time, marking a turning point in the relationship between Anthropic and the federal administration.

> The system became part of the daily analytical and planning flows within certain structures in the Department of Defense

From that point, the relationship between Anthropic and the Department of Defense lost the elements of institutional trust necessary for continued cooperation.

Anthropic did not enter cooperation with the Pentagon by accident. Since November 2024, the company, through a technical partnership with Palantir and Amazon Web Services, has implemented the Claude model on classified military networks with the DISA Impact Level 6 security standard, which represents one of the highest levels of permitted access in the US defence system.

At that time, Anthropic was the only commercial AI company with such status. In July 2025, the Department of Defense concluded a contract with Anthropic, with a maximum value of up to 200 million dollars, aimed at developing and applying artificial intelligence systems within defence operations.

By the end of 2025, Claude was integrated into the processes of intelligence analysis, operational planning, and simulation of combat scenarios.

At that stage, the system became part of the daily analytical and planning flows within certain structures in the Department of Defense. According to Pentagon officials, it was practically irreplaceable.

Classified as a supply chain security risk

According to Reuters, President Donald Trump's administration classified Anthropic as a "supply chain security risk," formally excluding it from federal use.

This decision was preceded by an ultimatum to the company's management: either Anthropic would remove the technical limitations built into the Claude model and permit its use for all legitimate Department of Defense purposes, or it would be treated as a security threat.

The disputed restrictions concerned a ban on mass surveillance of American citizens and a ban on the use of fully autonomous weapons systems without human supervision.

Anthropic CEO Dario Amodei rejected this request, stating that the company could not agree to conditions that exceeded the existing legal framework.

After the deadline, the decision was formalised on the same day. Federal agencies were ordered to suspend further use of Anthropic's technology, with a six-month transition period for systems in which Claude was already deeply integrated.

> Claude was used in intelligence analysis, target identification, and simulation of combat scenarios during Operation Epic Fury

The White House's public response was political, but the administrative measure was precise and legally structured, based on a mechanism that had until then been applied almost exclusively to foreign technology suppliers.

Less than a day later, the United States launched air and missile strikes against targets in Iran as part of Operation Epic Fury.

At that stage, Claude, through existing operational arrangements, remained active in certain structures of the United States Central Command (CENTCOM), where it was used in intelligence analysis, target identification, and simulation of combat scenarios.

The ban, therefore, took effect in political and contractual terms, but the system's operational dependence on already implemented technology prevented its immediate and complete enforcement.

Yes to OpenAI and no to Anthropic

In this context, the action taken by Sam Altman, CEO of OpenAI, on the same Friday night takes on a very different meaning from that presented in the media.

OpenAI announced that it had reached an agreement with the Pentagon to deploy its models on classified networks. The company presented the move to the public as a strategic success.

However, the crucial question is not who secured the contract but rather the circumstances surrounding its conclusion.

The conditions were the same as those that had previously led to Anthropic's exclusion.

OpenAI's agreement with the Pentagon includes a ban on mass surveillance of American citizens and a requirement for human accountability in the use of force, including autonomous weapons systems.

Altman publicly confirmed this, stating that these provisions are an integral part of the contract and that the Pentagon agreed to them, explaining that they are already contained in the current legal and political framework.

> OpenAI reached an agreement with the same institutions on terms that Anthropic had previously found unacceptable

The Pentagon, therefore, refused to accept from Anthropic what it accepted from OpenAI that same morning. And those were identical clauses.

After the Pentagon's decision was announced, Sam Altman described the treatment of Anthropic as a worrying precedent and expressed hope that the situation would be resolved differently.

At the same time, OpenAI reached an agreement with the same institutions on terms that Anthropic had previously found unacceptable.

Thus, the contract confirmed the framework already established by the exclusion decision. There was no public discussion about the compatibility of these positions.

Executive Order 12333 and the limits of AI oversight

One legal detail went almost unnoticed in public analyses. It concerns the way OpenAI's agreement with the Pentagon addresses Executive Order 12333, which governs the activities of the United States National Security Agency (NSA).

The contract stipulates that data collection and processing will comply with this order.

However, EO 12333 permits the interception of communications outside the United States, including those involving data of US citizens.

Anthropic specifically raised this issue during negotiations with the Department of Defense.

The company argued that the existing legal framework does not fully address the capabilities of modern artificial intelligence systems, particularly regarding the aggregation and analysis of large volumes of publicly available data for surveillance purposes.

Instead of insisting on additional contractual restrictions, OpenAI accepted reliance on the existing legal regime.

This reduced the issue that Anthropic considered a fundamental security problem to mere formal compliance with regulations not adapted to new technological possibilities.

> The Pentagon attempted to apply a mechanism developed to restrict foreign suppliers to a domestic company with a broad base of private users

A key consequence of these events concerns not the relationship between the federal government and the technology industry, but the future framework for legal and technical control of artificial intelligence in military systems.

Labelling Anthropic as a "supply chain risk" raised questions about the scope of this instrument.

According to the company's own legal interpretation, such a classification may apply exclusively to the contractual use of Claude within government structures, but not to its commercial users outside the federal administration.

In doing so, the Pentagon attempted to apply a mechanism developed to restrict foreign suppliers to a domestic company with a broad base of private users, representing a precedent without clear grounding in existing practice.

Political decisions, operational reality

The resulting legal dispute could raise questions about the limits of executive power in regulating commercial technology used for security purposes, as well as the relationship between contractual obligations and the broader market use of the same systems. These are issues that American courts have not previously addressed.

During the largest US military operation in decades, planning and analysis relied on a system whose manufacturer was designated a security risk on the same day

The outcome has already been reflected in operational practice. During the largest US military operation in decades, planning and analysis relied on a system whose manufacturer was designated a security risk on the same day.

The decision existed as an administrative act, but its implementation was not possible once operations were already underway.

This marks the boundary between political decisions and the reality of war, which was clear in this case.

It indicates how, in military structures that depend on new technologies and artificial intelligence systems, decisions can be made more quickly than they can be implemented, and how rules in practice are adapted to operational reality, rather than the other way around.