

February 28, 2026 Company

Our agreement with the Department of War

[Share](#)

Update on March 2, 2026

Throughout our discussions, the Department made clear it shares our commitment to ensuring our tools will not be used for domestic surveillance. To make our principles as clear as possible, we worked together to add additional language to our agreement.

This language makes explicit that our tools will not be used to conduct domestic surveillance of U.S. persons, including through the procurement or use of commercially acquired personal or identifiable information. The Department also affirmed that our services will not be used by Department of War intelligence agencies like the NSA. Any

We use cookies

We use cookies to help this site function, understand service usage, and support marketing efforts. Visit [Manage Cookies](#) to change preferences anytime. View our [Cookie Policy](#) for more info.

[Manage Cookies](#)[Reject non-essential](#)[Accept all](#)

OpenAI

including through the procurement or use of commercially acquired personal or identifiable information.

The Department of War plans to convene a working group made up of leaders from the frontier AI labs, cloud providers, and the Department's policy and operational communities. OpenAI will participate and expect this will be an important forum for ongoing dialogue on emerging AI capabilities, privacy, and national security challenges going forward.

These updates build on the framework we announced last week and we hope will help create a pathway for other labs to work with the Department going forward.

Yesterday we reached an agreement with the Pentagon for deploying advanced AI systems in classified environments, which we requested they also make available to all AI companies.

We think our agreement has more guardrails than any previous agreement for classified AI deployments, including Anthropic's. Here's why.

We have three main red lines that guide our work with the DoW, which are generally shared by several other frontier labs:

We use cookies

We use cookies to help this site function, understand service usage, and support marketing efforts. Visit [to change preferences anytime. View our Cookie Policy for more info.](#)

OpenAI

approach. We retain full discretion over our safety stack, we deploy via cloud, cleared OpenAI personnel are in the loop, and we have strong contractual protections. This is all in addition to the strong existing protections in U.S. law.

We believe strongly in democracy. Given the importance of this technology, we believe that the only good path forward requires deep collaboration between AI efforts and the democratic process. We also believe our technology is going to introduce new risks in the world, and we want the people defending the United States to have the best tools.

Our agreement includes:

1. Deployment architecture. This is a cloud-only deployment, with a safety stack that we run that includes these principles and others. We are not providing the DoW with “guardrails off” or non-safety trained models, nor are we deploying our models on edge devices (where there could be a possibility of usage for autonomous lethal weapons).

Our deployment architecture will enable us to independently verify that these red lines are not crossed, including running and updating classifiers.

2. Our contract. Here is the relevant language:

The Department of War may use the AI System for all lawful purposes, consistent with applicable law, operational requirements, and well-established safety and oversight protocols. The AI System will not be used to independently direct autonomous weapons in

We use cookies

We use cookies to help this site function, understand service usage, and support marketing efforts. Visit [to change preferences anytime. View our Cookie Policy for more info.](#)

OpenAI

system shall also not be used for domestic law-enforcement activities except as permitted by the Posse Comitatus Act and other applicable law.

3. AI expert involvement. We will have cleared forward-deployed OpenAI engineers helping the government, with cleared safety and alignment researchers in the loop.

FAQ

Why are you doing this?

First, we think the US military absolutely needs strong AI models to support their mission especially in the face of growing threats from potential adversaries who are increasingly integrating AI technologies into their systems. We originally did not jump into a contract for classified deployment, as we did not feel that our safeguards and systems were ready, and have been working hard to ensure that a classified deployment can happen with safeguards to ensure that red lines are not crossed.

We were—and remain—unwilling to remove key technical safeguards to enhance performance on national security work. That is not the correct approach to supporting the US military.

Second, we also wanted to de-escalate things between DoW and the US AI labs. A good future is going to require real and deep collaboration between the government and the AI

We use cookies

We use cookies to help this site function, understand service usage, and support marketing efforts. Visit [to change preferences anytime](#). View our [Cookie Policy](#) for more info.

OpenAI

responsible safeguards than earlier agreements, including Anthropic’s original contract. We think our red lines are more enforceable here because deployment is limited to cloud-only (not at the edge), keeps our safety stack working in the way we think is best, and keeps cleared OpenAI personnel in the loop.

We don’t know why Anthropic could not reach this deal, and we hope that they and more labs will consider it.

Do you think Anthropic should be designated as a “supply chain risk”?

No, and we have made our position on this clear to the government.

Will this deal enable the Department of War to use OpenAI models to power autonomous weapons?

No. Based on our safety stack, our cloud-only deployment, the contract language, and existing laws, regulation and policy, we are confident that this cannot happen. We will also have OpenAI personnel in the loop for additional assurance.

Will this deal enable the Department of War to use OpenAI models to conduct mass surveillance on U.S. persons?

No. Based on our safety stack, the contract language, and existing laws that heavily restrict DoW from domestic surveillance, we are confident that this cannot happen. We

We use cookies

We use cookies to help this site function, understand service usage, and support marketing efforts. Visit [to change preferences anytime. View our Cookie Policy for more info.](#)

OpenAI

As with any contract, we could terminate it if the counterparty violates the terms. We don't expect that to happen.

What if the government just changes the law or existing DoW policies?

Our contract explicitly references the surveillance and autonomous weapons laws and policies **as they exist today**, so that even if those laws or policies change in the future, use of our systems must still remain aligned with the current standards reflected in the agreement.

How do you address the arguments Anthropic made in their blog post about their discussion with the DoW?

In their post, Anthropic states two of their red lines (we have the same two red lines, plus a third: automated high-stakes decision making), and reasons they do not believe these red lines would be upheld in the contracts they had seen from the DoW at that time. Below is why we believe those same red lines would hold in our contract:

- **Mass domestic surveillance.** It was clear in our interaction that the DoW considers mass domestic surveillance illegal and was not planning to use it for this purpose. We ensured that the fact that it is not covered under lawful use was made explicit in our contract.
- **Fully autonomous weapons.** The cloud deployment surface covered in our contract

We use cookies

We use cookies to help this site function, understand service usage, and support marketing efforts. Visit [our Cookie Policy](#) to change preferences anytime. View our [Cookie Policy](#) for more info.

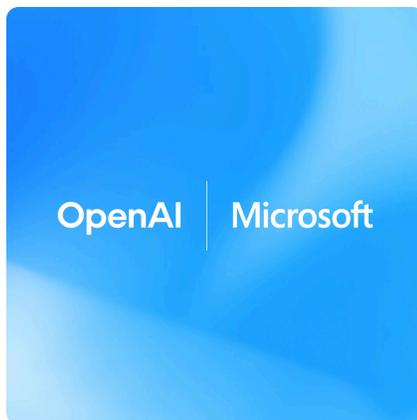
OpenAI

Author

[OpenAI](#)

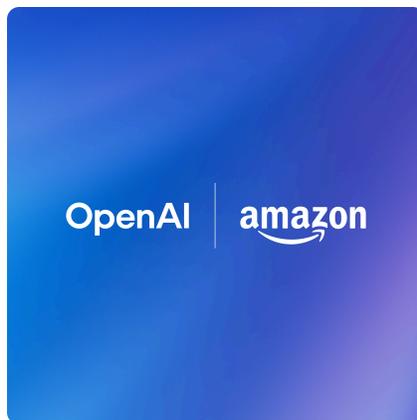
Keep reading

[View all](#)



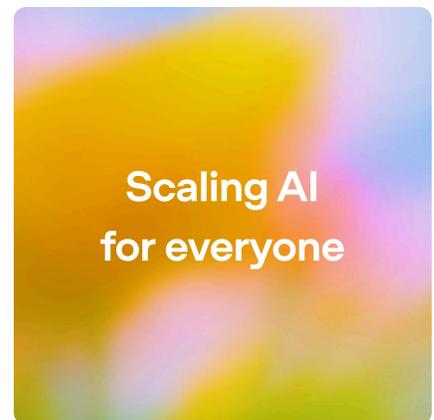
Joint Statement from
OpenAI and Microsoft

Company Feb 27, 2026



OpenAI and Amazon
announce strategic
partnership

Company Feb 27, 2026



Scaling AI for everyone

Company Feb 27, 2026

We use cookies

We use cookies to help this site function, understand service usage, and support marketing efforts. Visit [to change preferences anytime](#). View our [Cookie Policy](#) for more info.

OpenAI

Economic Research

Latest

Advancements

GPT-5.3 Instant

GPT-5.3-Codex

GPT-5

Codex

Safety

Safety Approach

Security & Privacy

Trust & Transparency

Pricing ↗

Download ↗

Sora

Sora Overview

Features

Pricing

Sora log in ↗

API Platform

Platform Overview

Pricing

API log in ↗

Documentation ↗

Developer Forum ↗

Company

About Us

Our Charter

Foundation

Careers

Brand

Support

Help Center ↗

More

News

Stories

Livestreams

Podcast

RSS



OpenAI © 2015–2026

English United States

We use cookies

We use cookies to help this site function, understand service usage, and support marketing efforts. Visit [our Cookie Policy](#) to change preferences anytime. View our [Cookie Policy](#) for more info.