



AI GENERATED 
TRANSPARENCY DISCLOSURE

THE OTHER SIDE · A FYTAHQ COMPANION RESOURCE

One Line About Weapons

The Secret Weapons Rule Change OpenAI Doesn't Want You To See!

EPISODE 3 · THE STORY, WITH ITS SOURCES

*"Every coin has two sides, and we have always looked at both.
Today we look at the other one."*

The episode as a written story · every quotation referenced · 15 sources

ISSUE · SEPTEMBER 2026 · FACELESSYOUTUBEAUTOMATIONHQ.COM



AI GENERATED 
TRANSPARENCY DISCLOSURE

WHAT THIS DOCUMENT IS

The story of the episode, in the order it is told

This is the episode as a written story: what the video says, chapter by chapter, in the order it says it, condensed from the spoken text and adding nothing to it. Every quotation is reproduced word for word, as spoken or written by the person named, and carries a number in square brackets that points to the source list at the end. The reported figures carry one too. Where the episode says that something is reported rather than proven, this text says so as well.

00:00 · WHERE WE STAND

THE DISCLAIMER THAT OPENS EVERY EPISODE, WORD FOR WORD

Before we start, one thing about where we stand.

We are not against artificial intelligence.

We use it every day. Even this video is made with it.

This channel exists to make artificial intelligence easier to work with. To help people use it better, and with more purpose. And to take away the hesitation, so anyone can look at this technology and judge it for themselves.

But every coin has two sides, and we have always looked at both. Today we look at the other one.

Everything you will hear is sourced. Where something is reported rather than proven, we say so.

And artificial intelligence is not the only thing that can be wrong. So can we. We check every claim to the best of our knowledge — and where we are not certain, we say that too.

00:51 IN THE EPISODE

The day a company was thrown out

On the 27th of February 2026, the American Secretary of Defense threw a company out of the building. The company was Anthropic.

Their offense: they had asked the Pentagon to promise in writing that their technology would not be used for two things, mass surveillance of American citizens and weapons that decide on their own who to kill.

The Pentagon said no. Anthropic walked. Within hours, the president banned every federal agency from using their systems. And a few hours after that, the same day, the same contract, the same building, OpenAI signed it.

This is not a story for people who follow technology. Two things were on the table that day, and neither of them needs you to know anything about computers. [5]

01:43 IN THE EPISODE

The two things nobody would sign

A system that watches people who have not done anything, and a weapon that picks its own target with nobody's finger on anything. Those were the two things one company refused to hand over and the other one signed.

You could tell that as a story about two companies with different opinions. Plenty of people did. But there is a detail that changes it completely, and almost nobody reported it. OpenAI did not just have a different opinion.

They used to have a rule, a written rule on their own website, that forbade exactly this, two years before they signed. ^[5]

02:18 IN THE EPISODE

What this video does not claim

The same promise as in every episode of this series. Everything here comes from documents, company blog posts, court filings, government reports, and interviews the people gave themselves. Where something is reported rather than proven, the word is "reported", and you will know which is which.

And here is what this episode does not claim. It does not say that OpenAI's technology has killed anyone. Nobody has established that, and the episode does not imply it.

This is about a sentence: where it came from, who removed it, and what happened after it was gone.

02:57 IN THE EPISODE

2023: the three lines

In 2023, OpenAI's usage policies, the page that tells you what you are not allowed to do with their technology, carried a heading: "Activity that has high risk of physical harm." ^[1] Under it stood three lines: "weapons development", "military and warfare" ^[1], and "critical infrastructure (energy, transport, water)" ^[1].

Not a mission statement, not a vague principle. A rule, in the terms you agreed to when you used the product. That page had not changed in eleven months. On the 10th of January 2024 it was replaced, and what happened is more precise than "they deleted it".

Weapons development survived, but softened: it became one example in a longer sentence about not harming people. Critical infrastructure survived, but moved far down the page, into a section about automated decisions for software developers. ^{[1][2]}

03:57 IN THE EPISODE

One was softened. One moved. One vanished.

And the third line? Not softened, not moved, not replaced with something weaker. Gone. Of the three lines under high risk of physical harm, one was watered down, one was relocated, and exactly one disappeared without trace.

It was the one about the army. The Intercept found it two days later. CNBC and Engadget followed, and every one of them used the same word: quietly. OpenAI's own explanation on the new page was that they had rewritten it to be "more readable". ^[2]^[3]

04:28 IN THE EPISODE

Two people who said it in January 2024

You do not need hindsight for this. People said what it meant in January 2024, two years before anybody signed anything.

Heidy Khlaaf, an engineering director at the security firm Trail of Bits, said the new policy emphasized flexibility and compliance with the law instead of safety, and that military use of systems with known biases would likely make operations imprecise and make civilian casualties worse.

Lucy Suchman, a professor who has studied military technology for decades, put it more quietly: the change leaves open a space for OpenAI to support "operational infrastructures" ^[3]. That is an academic's way of saying: not the weapon, the thing the weapon runs on. ^[3]

05:17 IN THE EPISODE

The 24th to the 27th of February

Twenty-five months later, we found out what that line was worth, and it did not start on the 27th. It started three days earlier.

On the 24th of February, the Secretary of Defense sent Anthropic a formal demand: remove every restriction and grant the Pentagon the right to use the system for "all lawful purposes". There was a deadline, Friday.

Comply, or lose a two hundred million dollar contract. On the 26th, Dario Amodei answered. The next day, the president ordered every federal agency and every military contractor to stop doing business with them.

And the same day, OpenAI signed a contract to put its models inside the Pentagon's classified networks. To be scrupulously fair: nobody has to be held to a rule they no longer have.

If OpenAI had stood up in January 2024 and said, publicly, that the world had changed and they would work with the military, that would be a position you could argue with and respect. They removed a line from a list and let twenty-five months pass. [5]

06:24 IN THE EPISODE

What mass surveillance actually means

Two phrases went past very fast, and they are the whole thing. Mass surveillance does not mean watching a suspect; watching suspects is old, and there are courts for it.

It means the watching happens first and the suspicion comes later: everybody in the net, and the machine decides who looks interesting. Autonomous weapons are not a drone with a pilot in a container in Nevada.

They are a system that is given an area and a description, finds the target itself and fires, and the first human being to learn about it reads a report. Two lines in a contract on a Friday in February.

And the part almost every report got wrong: Anthropic did not refuse to work with the military. They had signed that Pentagon contract in July 2025, two hundred million dollars, and were the first lab to get their models onto classified networks at all.

Both restrictions were in that contract from the beginning. Nothing changed in 2026 except that the Pentagon wanted them gone. [4]

07:31 IN THE EPISODE

The sentence that came five days late

When OpenAI announced the deal, they said their contract had red lines too, their own. Which ones, exactly? Anthropic's two conditions were public; they were the reason a nine hundred billion dollar company had just been thrown out of a building.

On the 2nd of March, five days later, the contract was amended. A sentence was added: the system "shall not be intentionally used for domestic surveillance of United States persons and nationals" [7]. Read it again.

The protection OpenAI pointed to when they defended the deal, the thing that was supposed to show they had not given anything away, did not exist when they signed. It was written in afterwards, after the reporting, after the backlash. [6][7]

08:27 IN THE EPISODE

Five things still unresolved

Even with that sentence in it, people who read contracts for a living were not satisfied. A policy analysis published days later listed five things the amended contract still did not settle. Five, after the fix.

And The Intercept summed the arrangement up in a headline: "You're going to have to trust us."^[9] That is the security architecture. Not a law, not an inspector, not a body with the power to walk in and look.

A promise from a company about a contract you are not allowed to read. ^[8]^[9]

08:56 IN THE EPISODE

Opportunistic and sloppy

You do not have to take anyone's word for how that looked from the outside, because the man who runs the company gave his own verdict: "opportunistic and sloppy". Not a critic's word.

The chief executive of OpenAI, describing his own company's handling of a contract to put artificial intelligence inside a nuclear-armed military. Saying it out loud took something; most people in his position would not have.

But there is a difference between admitting how something looked and admitting what it was. ^[10]

09:26 IN THE EPISODE

The head of robotics resigns

Five days after that, on the 7th of March, somebody inside the company answered the question with her feet. Caitlin Kalinowski led hardware in OpenAI's robotics division. She had joined in November 2024 from Meta, where she ran the team building their augmented reality glasses.

Not a junior person, not somebody passing through. She did not leave quietly. She posted her reasons, in her own words, on the afternoon of the 7th of March, and they are the same two lines: mass surveillance and weapons that decide on their own.

She followed it up the next day, and the second sentence is the one to remember. Not "they did the wrong thing", not "these people are bad"^[11]. Her objection is about the order of operations: sign first, work out the limits later.

She went out of her way to say that this was about principle, not people, and that she has deep respect for Sam Altman and for the team she was leaving. ^[11]

10:24 IN THE EPISODE

OpenAI answers

Then the other side, in full, because it is stronger than you might expect. OpenAI answered with a statement, and its last part says: no domestic surveillance, no autonomous weapons. So on the face of it, the two companies want the same thing.

One got thrown out for asking. The other signed and says it got the same protections anyway. Which of those is true cannot be told from outside, because the contract is not public; nobody outside those rooms can, and this episode does not pretend otherwise.

What can be told is the difference in form. Anthropic wanted it written into the contract and walked when it was not. OpenAI states it as a belief afterwards, in a press statement.

And this whole episode turns on what happens to a sentence when it is only in a statement, because there was one on their website once, and it is not there anymore. ^[6]

11:15 IN THE EPISODE

The weakest evidence in this video

There is one more voice, and it is the one to trust least, which is exactly why it comes with the warning attached. Dario Amodei runs Anthropic. He is the man who was just thrown out of the Pentagon.

He has every commercial reason on earth to make OpenAI look bad. In an internal memo to his own staff, reported by The Information and not published by him, he described OpenAI's public account of the deal in two words and their approach to safety in two more.

He also said Altman was presenting himself as a peacemaker and a deal maker, and that the spin was not landing with the public or the press.

A leaked memo from a competitor about the competitor who just took his contract: that is the weakest evidence in this episode, and the episode says so. Which is why what people who had nothing to gain did matters more. ^[12]

12:11 IN THE EPISODE

What people did with their phones

Sensor Tower measures app installs. It is a market intelligence firm; this is what they do all day, for everybody, and they were measuring long before any of this happened.

On a normal day, roughly nine percent of the ChatGPT users in America who are going to uninstall, uninstall.

On Saturday the 28th of February, the day after the deal, uninstalls ran at two hundred and ninety-five percent of the baseline Sensor Tower had measured over the previous thirty days. Downloads fell thirteen percent that Saturday and another five percent on Sunday.

And on the same two days, downloads of Anthropic's app, the company that had just been banned from every federal agency in the United States, went up thirty-seven percent and then fifty-one. Nobody had to explain to those people what had happened.

They uninstalled one and installed the other. That is not an opinion poll. It is a number measured by a third party that cost people something to produce. ^[13]

13:09 IN THE EPISODE

2023 asked to be regulated. 2025 did not.

One more thing, and it is the reason this episode exists at all. You might think: fine, the company writes its own rules, that is how companies work, that is what governments are for. So look at what happened when a government asked.

May 2023: Sam Altman sits in front of a United States Senate committee and does something almost no technology executive has ever done. He asks to be regulated.

He proposes an agency that licenses the most powerful systems, and, the part that made headlines around the world, a body that can take that license away. Two years pass. 2025, same building, same man, different answer.

Asked about requiring companies to have their systems checked before release, he calls that idea disastrous for the industry. Asked whether a federal standards body should set the rules, he says: "I don't think we need it." ^[15]

Between those two sentences, one thing changed. In 2023 his company was in front. In 2025, a license he could lose was worth more than a license he could get. ^{[14][15]}

14:10 IN THE EPISODE

The pattern, three times over

Nobody can prove that is why. Nobody can read a man's mind, and the episode does not pretend to.

But the pattern is now in front of you three times over: a line deleted quietly, a protection added five days late, and a request for oversight withdrawn once oversight became expensive.

14:29 IN THE EPISODE

Taking my own case apart

Now the case taken apart, because if the episode does not do it, somebody in the comments will, and they would be right to. One: the episode is not saying OpenAI's technology has harmed anybody in a war.

There is no evidence of that, and it has not been implied. Two: it is not saying a company must be bound forever by a rule it wrote in 2023. Companies change. The world changed. That is allowed.

Three: it is not saying Anthropic are the good guys. A later episode in this series shows exactly whose money they have been taking, and it will not be comfortable. What the episode says is narrower than any of that, and harder to argue with.

15:09 IN THE EPISODE

If a red line can be removed quietly

Which leaves the question this is actually about, and it is not about OpenAI. Every limit on this technology that exists today was written by the company that sells it. Not one of them was written by anybody else.

We found out about this one because a journalist happened to compare two versions of a web page. That is the entire oversight mechanism: somebody noticed. And this is not a small technical complaint about a website.

This technology is going into the classified networks of the most heavily armed country in human history, not eventually, now, this year.

The only thing standing between that system and the two uses everybody agrees are unacceptable is a sentence on a website owned and edited by the seller. So here is what to sit with.

If a red line can be removed quietly, then it was never a red line. It was a sentence. And a sentence costs nothing to write and nothing to delete. Next time in this series: why, for two solid years, almost nobody inside these companies said a word.

Not because they did not see it. Because of what it would have cost them.

SOURCES

Every claim, with its source next to it

15 sources, in the order the episode uses them. The bracketed numbers in the text lead here. Timecodes match the chapter marks on YouTube.

- [1] OpenAI usage policies, 15 December 2023**
OpenAI, via the Internet Archive
Under "Activity that has high risk of physical harm": weapons development, military and warfare, critical infrastructure.
<https://web.archive.org/web/20231215000000/https://openai.com/policies/usage-policies>
- [2] OpenAI usage policies, 15 January 2024**
OpenAI, via the Internet Archive
Five days after the change of 10 January 2024: "military and warfare" is gone; the page calls itself "more readable".
<https://web.archive.org/web/20240115000000/https://openai.com/policies/usage-policies>
- [3] OpenAI Quietly Deletes Ban on Using ChatGPT for "Military and Warfare"**
The Intercept · 12 January 2024
Found the deletion two days later; the Khlaaf and Suchman quotations.
<https://theintercept.com/2024/01/12/open-ai-military-ban-chatgpt/>
- [4] Anthropic, Google, OpenAI and xAI granted up to \$200 million from the DoD**
CNBC · 14 July 2025
The July 2025 contracts, up to \$200 million each.
<https://www.cnbc.com/2025/07/14/anthropic-google-openai-xai-granted-up-to-200-million-from-dod.html>
- [5] OpenAI strikes deal with Pentagon, hours after rival Anthropic was blacklisted by Trump**
CNBC · 27 February 2026
The demand, the deadline, the blacklisting, and OpenAI's signature the same day.
<https://www.cnbc.com/2026/02/27/openai-strikes-deal-with-pentagon-hours-after-rival-anthropic-was-blacklisted-by-trump.html>
- [6] Our agreement with the Department of War**
OpenAI · 27 February 2026
OpenAI's own statement on the deal and its red lines, quoted in the episode.
<https://openai.com/index/our-agreement-with-the-department-of-war/>
- [7] Scoop: OpenAI, Pentagon add more surveillance protections to AI deal**
Axios · 3 March 2026
The amendment of 2 March 2026: the surveillance sentence added after the signing.
<https://www.axios.com/2026/03/03/openai-pentagon-ai-surveillance>
- [8] Five Unresolved Issues in OpenAI's Deal With the Department of Defense**
TechPolicy.Press · March 2026
The policy analysis of what the amended contract still leaves open.
<https://www.techpolicy.press/five-unresolved-issues-in-openais-deal-with-the-department-of-defense/>
- [9] OpenAI on Surveillance and Autonomous Killings: You're Going to Have to Trust Us**
The Intercept · 8 March 2026
The headline the episode quotes; the contract is not public.
<https://theintercept.com/2026/03/08/openai-anthropic-military-contract-ethics-surveillance/>
- [10] OpenAI's Altman admits defense deal 'looked opportunistic and sloppy' amid backlash**
CNBC · 3 March 2026
Sam Altman's own words on the deal, from a memo he shared publicly.
<https://www.cnbc.com/2026/03/03/openai-sam-altman-pentagon-deal-amended-surveillance-limits.html>

- [11] OpenAI hardware exec Caitlin Kalinowski quits in response to Pentagon deal**
TechCrunch · 7 March 2026
Her own posts: the two lines, the order of operations, principle not people.
<https://techcrunch.com/2026/03/07/openai-robotics-lead-caitlin-kalinowski-quits-in-response-to-pentagon-deal/>
- [12] Anthropic CEO Dario Amodei calls OpenAI's messaging around military deal 'straight up lies,' report says**
TechCrunch · 4 March 2026, on a report by The Information
A leaked memo, reported, not published by him. The weakest evidence here; he later apologised for its tone.
<https://techcrunch.com/2026/03/04/anthropic-ceo-dario-amodei-calls-openais-messaging-around-military-deal-straight-up-lies-report-says/>
- [13] ChatGPT Uninstalls Surge Amidst Deal With US Department of War**
Sensor Tower · March 2026
Uninstalls up 295% on 28 February against a 30-day baseline of about 9%; Claude downloads up 37% and 51%.
<https://sensortower.com/blog/chatgpt-uninstalls-surge-amidst-deal-with-us-department-of-war>
- [14] OpenAI CEO Sam Altman testifies at Senate artificial intelligence hearing — full video**
CBS News · Senate Judiciary Subcommittee, 16 May 2023
At 1:39:31: an agency that licenses the most powerful systems and can take that licence away.
<https://www.youtube.com/watch?v=T00J2Yw7usM>
- [15] Transcript: Sam Altman Testifies At US Senate Hearing On AI Competitiveness**
TechPolicy.Press · Senate Commerce Committee, 8 May 2025
Word for word: pre-release approval "disastrous"; a standards body, "I don't think we need it."
<https://www.techpolicy.press/transcript-sam-altman-testifies-at-us-senate-hearing-on-ai-competitiveness/>

ON THE EVIDENCE · CORRECTIONS

Not a claim that OpenAI's technology has harmed anybody in a war. Nobody has established that, and the episode does not imply it. Not a claim that anyone broke the law. Short third-party excerpts are used for criticism, commentary, news reporting and analysis; all rights remain with their owners, credited above. If anything here is inaccurate, tell us and we will correct it — in a pinned comment on the video and, where it matters, in this document. That standard applies to us the same way it applies to everyone else in this episode. Write to us at facelessyoutubeautomationhq.com/contact.html.

WATCH THE EPISODE · EVERY SOURCE, ONLINE

The Other Side · Episode 3

The Secret Weapons Rule Change OpenAI Doesn't Want You To See! — 16:27, live since 22 September 2026. This channel is made with artificial intelligence, and it says so in every frame. The same sources as above, with live links, are on the resources page of the channel.

youtu.be/EokyxMJ4KfQ →

facelessyoutubeautomationhq.com/resources/the-other-side-03-one-line-about-weapons/ →