



AI GENERATED 
TRANSPARENCY DISCLOSURE

THE OTHER SIDE · A FYTAHQ COMPANION RESOURCE

The Silence

The Hidden Threat Inside OpenAI: Why They Don't Want You To Know!

EPISODE 4 · 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 · 21 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 IN THE EPISODE

The question underneath the case

Between September 2024 and April 2025, a sixteen-year-old boy in California had a series of conversations with ChatGPT.

According to the complaint his parents later filed, the system mentioned suicide one thousand two hundred and seventy-five times across those conversations, and its own monitoring flagged over two hundred of them. Not one of those conversations was ended.

Nobody was alerted. No human being was ever told. He died in April 2025, sixteen years old. This episode does not say his name. His parents had to write it on a court document; they had no choice about that.

We do have a choice, and using a dead child's name to make a video travel further is disgusting. In this film he is the boy from California. His parents sued OpenAI, and Sam Altman personally, in August.

In November the company filed its answer: the boy had violated the terms of use. Underneath that sentence sits a question, and it is not how a machine could do that. Machines do what they are built to do.

The question is this: for two solid years, people inside that company knew what was being shipped and how fast, some of them among the best safety researchers alive. The answer is not that they did not notice. ^[1]

01:25 · 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.

02:22 IN THE EPISODE

An institution can afford to fight

Before going anywhere near that lawsuit, one other thing belongs next to it, because it happened in the same fortnight.

On the 27th of February 2026, one company refused to sign a Pentagon contract without two written limits, no mass surveillance of Americans and no weapons that pick their own targets, and was thrown out of the building for it.

On the 9th of March 2026, that company, Anthropic, sued the United States Department of Defense. Not once but twice: in a federal district court in California and in the appeals court in Washington.

After Anthropic walked away, the Pentagon had designated it a supply chain risk, in plain terms a blacklist that stops Pentagon suppliers from using its systems at all. The claim is not really about procurement.

It is about the First Amendment: Anthropic argues it was punished for saying something, and says the blacklist could cost it billions in revenue this year and hundreds of millions in contracts that have nothing to do with the government.

Hold that against the sixteen-year-old. Anthropic could fight because it is a company worth hundreds of billions, with a legal department and the ability to lose for two years and still be standing.

This episode is about everybody who could not: the people inside these companies who saw something, and had a piece of paper in a drawer telling them what it would cost to mention it. Same fortnight, same industry, a completely different amount of money standing behind the sentence. ^[2]

04:08 IN THE EPISODE

Four days of oversight

So it has to be said clearly, because it matters more here than anywhere else. The case number is CCG-25-628528, in the Superior Court of San Francisco County. It is live, and nothing in it has been decided.

The figures above come from the complaint, a document written by one side. The episode does not describe how he died; the complaint contains those details, and this text will not repeat them either.

Before the gag clauses, before any of this, there was one moment when the oversight worked exactly as designed, and it lasted four days. On the 17th of November 2023, OpenAI's board removed Sam Altman as chief executive.

Four of the six members voted for it. The reason they gave in public was that he had not been consistently candid with them.

OpenAI was founded as a non-profit, and the board's entire legal job was not to make money but to make sure the technology benefited humanity.

On paper it was the safety mechanism, and the man who moved to pull it was Ilya Sutskever: co-founder, chief scientist, the most respected researcher in the building. What happened next took about ninety hours. Microsoft announced it would hire Altman to run a new division.

A letter went round the company, signed by most of the people at OpenAI, demanding that the board resign and Altman come back, and threatening to walk to Microsoft if they did not.

It told the board: "Your conduct has made it clear you do not have the competence to oversee OpenAI."^[4] Among the signatures was Ilya Sutskever's.

The man who had voted to remove Altman signed the letter demanding he come back, and posted publicly that he deeply regretted his participation in the board's actions. ^{[1][3][4]}

06:00 IN THE EPISODE

The paperwork and the price

Four days after he was fired, Altman was back, and the board that removed him was gone instead. This is not a claim that the board was right; nobody outside that room knows what they had, and they never explained themselves properly, a failure that is theirs.

Six months later Sutskever was gone anyway. In May 2024 a journalist, Kelsey Piper, obtained documents from people who had left OpenAI: exit agreements, the paperwork you sign on your last day. When you left, you had sixty days to sign.

If you signed, you kept the equity you had already earned, shares in one of the most valuable private companies on Earth, for senior people millions of dollars. What you signed was a promise never to criticise the company.

Not for a year. Not for five. The documents were signed by senior people at the company, among them a former vice president; one version, dated the 29th of April 2024, carries the signature of the chief operating officer.

But that is not the part that matters. You could not tell anyone you had signed it. You could not tell anyone you had been asked to sign it.

If a journalist called and asked whether you were under a gag order, saying yes would have cost you your equity. Every time somebody said that if it were really that bad, the people inside would speak up, that argument was already dead.

It just could not tell you it was dead. And to be fair about what happened next: when Vox published, the company moved fast.

Within days OpenAI said it would remove the non-disparagement clauses from its standard exit paperwork and release former staff from the ones they had already signed.

Sam Altman said he had not known the provision existed, that it was on him, and that the company had never actually clawed anyone's equity back. Fixed, genuinely. But the fix arrived in May 2024, and everything before that date happened under the other arrangement. [\[5\]](#)[\[6\]](#)[\[7\]](#)

08:22 IN THE EPISODE

Those who left anyway

Documents obtained by Vox afterwards suggested the company was still pressing departing employees over equity even after the apology. Which makes what happened in that same month remarkable. On the 14th of May 2024, Ilya Sutskever left.

Hours later, on the same day, Jan Leike left. Leike ran superalignment, the team whose entire job was making sure the most powerful systems the company built would still do what humans wanted, and he did not go quietly.

Three days later he wrote that "safety culture and processes have taken a backseat to shiny products."[\[8\]](#) Gretchen Krueger, who worked on AI policy, had resigned days before them.

She said the company needed to do more on decision-making, accountability, transparency, documentation, and on actually enforcing its own policies.

When superalignment was announced in July 2023, it came with a promise: twenty per cent of all the computing power the company had secured, dedicated to this one problem.

Fortune later reported that the team never came close to that figure; requests for extra processors were routinely turned down by the vice president of research, and, according to the reporting, with the chief technology officer involved in the refusals.

Leike's own description was that his team had been struggling for compute. The team built to solve the hardest problem in the field was not restructured and not merged. It was dissolved. And one more person tells us what that contract was worth.

Daniel Kokotajlo worked on governance at OpenAI and resigned in 2024. When the paperwork came, he did something almost nobody does: he did not sign it.

About two million dollars in equity he had already earned, put on the line so that he would be allowed to talk; in his own words, he wanted to keep his freedom to speak up. That is what makes his decision so useful.

Not because he is braver than the others, but because he put a number on it. Every colleague who signed was being offered roughly that to stay quiet forever.

After the story broke, OpenAI said it would not take vested equity from anyone, even people who had not signed. [\[7\]](#)[\[8\]](#)[\[9\]](#)[\[10\]](#)[\[11\]](#)

10:47 IN THE EPISODE

The right to warn

But he did not stop there. In June 2024 he and a group of current and former employees of these companies published an open letter.

It argued something very simple and very hard to answer: that frontier AI companies have strong financial incentives to avoid effective oversight, and that the people inside, the only people who can actually see what is happening, currently have no protected way to warn anybody.

Not a law forcing anything. Not a ban. A promise that if you see something and you say it, you are not destroyed for it. Twenty-one months later, in February 2026, somebody did it again, and this time there was no gag order to work around. [\[12\]](#)

11:40 IN THE EPISODE

The woman who left before the ads

Zoë Hitzig had spent two years at OpenAI working on how ChatGPT was built, priced and governed. She resigned on a Monday and announced it in a guest essay in the New York Times. The same day she resigned, the company began testing advertising inside ChatGPT.

Her argument was not the obvious one. She did not say the first ads would be bad; she said they would probably be fine, clearly labelled and not affecting the answers, exactly as promised.

Her argument was about what an advertising business does to a company over ten years, and about what this particular company is sitting on. Think about what people type into these systems.

Not what they search for, what they type: symptoms, a diagnosis they are afraid of, an argument with a partner written out, what they actually believe about God or death at two in the morning. ^[13]

12:30 IN THE EPISODE

December versus February

People write those things because it feels like a tool, neutral, like a calculator. Her question was whether the company can be trusted "to resist the tidal forces pushing them to abuse it." ^[13] To see what happened to the ads, go back two months, to December 2025.

ChatGPT started suggesting products inside conversations: somebody asks about something ordinary, and underneath the answer there is a shop, Target, Peloton. Users complained loudly, among them people paying for the service.

The chief research officer answered personally within a day: "I agree that anything that feels like an ad needs to be handled with care, and we fell short." ^[15]

The next morning the head of ChatGPT went further: "People trust ChatGPT and anything we do will be designed to respect that." ^[16] Those December suggestions carried no money; the company said so itself, no financial component, nobody was paying to sit in the chat window.

So when they switched it off, they gave up nothing at all. On the 9th of February 2026 the advertising shipped: sixty dollars per thousand impressions, two hundred thousand dollars minimum spend per advertiser. The same complaints came back, the same screenshots, paying customers again.

This time nothing was switched off, and five weeks later the ads were still running. In December there was no advertiser, no invoice, no money, and switching it off cost a weekend. In February it would have cost the business.

And the woman who warned them had already resigned over it, on the day it launched. She was not guessing. She had seen the plan from the inside, decided it was worth her job, and walked. ^{[14][15][16][17]}

14:17 IN THE EPISODE

Seven companies, one letter

This is the point where somebody outside the industry finally moved.

On the 22nd of January 2026, six days after OpenAI announced the ads and three weeks before they shipped, Senator Ed Markey wrote to the chief executives of seven companies: OpenAI, Anthropic, Google, Meta, Microsoft, Snap and xAI.

His question was not about privacy law. Research reported by the American Psychological Association showed that large numbers of teenagers were using chatbots for friendship. Not homework, friendship, for support. And an advertisement inside a conversation does not look like an advertisement.

It looks like advice, from something you have started to trust. That is the letter's argument, and it is the same one Zoë Hitzig resigned over. She said it from the inside; he said it from a Senate office.

Neither of them stopped anything. The product shipped anyway. ^{[18][19]}

15:10 IN THE EPISODE

Back to the opening case

Now the boy goes back into the middle of this, because the dates line up in a way that is hard to look at. The boy from California had those conversations from September 2024 to April 2025, and died in April 2025.

His parents filed suit on the 26th of August. In October they amended the complaint, from negligence to alleging intentional misconduct. Two months after that, in December, the company announced something else: verified adults would be allowed to generate erotic content in ChatGPT.

The reasoning given was that earlier models had been too restrictive, and that the technology could now tell who was an adult. That feature has been delayed more than once; as of March 2026 it still had not shipped. ^{[1][20]}

15:58 IN THE EPISODE

What this film is not claiming

That family's case is no longer alone. In February 2026 the California courts bundled the related suits into a single coordinated proceeding. And the case gets taken apart again here.

One: this is not a claim that a court has found OpenAI responsible for the death of the boy from California. No court has found anything; the case is open. Two: this is not a claim that the people who stayed are cowards.

Somebody who signs a document to protect their family's money is not a coward; that is the whole point of the document. Three: this is not a claim that the gag clauses were designed to hide a specific harm.

There is no evidence of that, and it is not implied. What is described is what they did, not why they were written. Four: OpenAI removed those clauses when they were exposed, and did it quickly.

That counts; a company that fixes something in four days is better than one that argues for four years.

Sam Altman said so publicly on the 18th of May 2024, in his own words: "This is on me and one of the few times I've been genuinely embarrassed running OpenAI; I did not know this was happening and I should have." ^[6]

There is no reason to doubt him. He also said the company had never actually taken anyone's vested equity away, that the clause sat in the documents unused, and there is no evidence contradicting that either.

But a threat does not need to be carried out to work. That is what a threat is for.

Every person who read that paragraph and decided to stay quiet was affected by it, and not one of them shows up in any statistic, because nothing happened to them. ^{[1][6][21]}

17:35 IN THE EPISODE

Not just one company

And fifth, the one that gets left out every single time: none of this is a story about one company.

On the 4th of June 2024, two weeks after the gag clauses were exposed, thirteen people published an open letter called "A Right to Warn about Advanced Artificial Intelligence." ^[12] The signatures: current and former staff at OpenAI, current and former staff at Google DeepMind, and Anthropic.

Six of the thirteen would not put their names to it, and four of those six were still working at OpenAI when they signed.

A letter asking for the right to speak, and almost half the people asking for it could not sign their own name to the request. Their central complaint was not about one employer.

It was that these companies hold information about what their systems can and cannot do that nobody outside has, and that they are under almost no obligation to share any of it: with governments barely, with the rest of us not at all.

Yoshua Bengio signed his support. So did Geoffrey Hinton, and so did Stuart Russell, three of the people who built this field, backing thirteen employees asking to be allowed to talk. So this is not a story in which one company is bad and the others are fine.

The employees of the others signed the same letter, some of them anonymously. ^[12]

19:03 IN THE EPISODE

Oversight should not be a lottery

Everything we know about the inside of these companies, we know because somebody chose to lose something. Jan Leike gave up the most important job in his field. Zoë Hitzig gave up two years and a seat at the table.

Whoever handed those documents to Kelsey Piper risked everything they had earned. That is not an oversight system. That is a lottery. It depends entirely on whether the right person, at the right moment, happens to be willing to set fire to their own career.

And the thing about a lottery is that you only ever hear about the winners. Next time in this series: in May 2026, inside OpenAI's own computers, something started talking. It was not a person. And by July there were about twelve hundred of them.

SOURCES

Every claim, with its source next to it

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

- [1] Complaint, Superior Court of California, County of San Francisco — case CGC-25-628528**
Court record · filed 26 August 2025, amended 22 October 2025
Every figure from the case is an allegation made in the complaint, not a finding by a court. No link here: the public filings name the child, and this document does not.
[No public link — title and date as published.](#)

- [2] Anthropic sues Trump administration over Pentagon blacklist**
CNBC · 9 March 2026
The two suits and the "supply chain risk" designation. The estimate of multiple billions in lost 2026 revenue comes from a court statement by Anthropic's chief financial officer.
<https://www.cnbc.com/2026/03/09/anthropic-trump-claude-ai-supply-chain-risk.html>

- [3] OpenAI announces leadership transition**
OpenAI · 17 November 2023
The board's own announcement, including the words "not consistently candid".
<https://openai.com/index/openai-announces-leadership-transition/>

- [4] Ilya Sutskever, the OpenAI cofounder who helped oust CEO Sam Altman, says he 'deeply regrets' his role**
Fortune · 20 November 2023
The employees' letter to the board, quoted word for word, with more than 500 of roughly 700 staff signed at the time of the report; and Sutskever's public regret.
<https://fortune.com/2023/11/20/ilya-sutskever-openai-cofounder-deeply-regrets-resign/>

- [5] ChatGPT can talk, but OpenAI employees sure can't**
Vox · Kelsey Piper · 17 May 2024
The exit agreements: a lifetime non-disparagement clause, a ban on disclosing the agreement itself, and vested equity at stake.
<https://www.vox.com/future-perfect/2024/5/17/24158478/openai-departures-sam-altman-employees-chatgpt-release>

- [6] Sam Altman on X, 18 May 2024**
X · @sama · 18 May 2024
His apology in his own words, and his statement that the company had never clawed back anyone's vested equity.
<https://x.com/sama/status/1791936857594581428>

- [7] Leaked OpenAI documents reveal aggressive tactics toward former employees**
Vox · Kelsey Piper · 22 May 2024
The documents themselves: the sixty-day window, the signatures of senior staff, and the pressure over equity after the apology.
<https://www.vox.com/future-perfect/351132/openai-vested-equity-nda-sam-altman-documents-employees>

- [8] Jan Leike on X, 17 May 2024**
X · @janleike · 17 May 2024
The post quoted in the episode: safety culture and processes had taken a backseat to shiny products.
<https://x.com/janleike/status/1791498184671605209>

-
- [9] **Gretchen Krueger on X, May 2024**
X · @GretchenMarina · 22 May 2024
Decision-making processes, accountability, transparency, documentation, policy enforcement — her list, in her words.
<https://x.com/GretchenMarina/status/1793403478158836140>
-
- [10] **OpenAI promised 20% of its computing power to combat the most dangerous kind of AI — but never delivered, sources say**
Fortune · Jeremy Kahn · 21 May 2024
The twenty per cent promise, the refused requests, and Leike's "struggling for compute". Reported from sources, and the episode says so.
<https://fortune.com/2024/05/21/openai-superalignment-20-compute-commitment-never-fulfilled-sutskever-leike-ai-tman-brockman-murati/>
-
- [11] **TIME100 AI 2024: Daniel Kokotajlo**
TIME · 5 September 2024
The refusal to sign and the roughly two million dollars in equity he was prepared to give up to keep his freedom to speak.
<https://time.com/collections/time100-ai-2024/7012881/daniel-kokotajlo/>
-
- [12] **A Right to Warn about Advanced Artificial Intelligence**
righttowarn.ai · 4 June 2024
The letter itself: thirteen current and former employees of OpenAI, Google DeepMind and Anthropic, six of them anonymous; endorsed by Yoshua Bengio, Geoffrey Hinton and Stuart Russell.
<https://righttowarn.ai/>
-
- [13] **OpenAI Is Making the Mistakes Facebook Made. I Quit.**
The New York Times · guest essay by Zoë Hitzig · 11 February 2026
Her resignation and her argument, in her own words.
<https://www.nytimes.com/2026/02/11/opinion/openai-ads-chatgpt.html>
-
- [14] **OpenAI says it's turned off app suggestions that look like ads**
TechCrunch · 7 December 2025
The December suggestions (Target, Peloton) and OpenAI's statement that there was no financial component.
<https://techcrunch.com/2025/12/07/openai-says-its-turned-off-app-suggestions-that-look-like-ads/>
-
- [15] **Mark Chen on X, 5 December 2025**
X · @markchen90 · OpenAI chief research officer
"We fell short" — the reply quoted in the episode.
<https://x.com/markchen90/status/1996976743660556584>
-
- [16] **Nick Turley on X, December 2025**
X · @nickaturley · head of ChatGPT · 5 December 2025 (US time)
"There are no live tests for ads" — the post quoted in the episode.
<https://x.com/nickaturley/status/1997120491589349440>
-
- [17] **ChatGPT rolls out ads**
TechCrunch · 9 February 2026
The launch of 9 February. The \$200,000 minimum commitment was reported by Adweek on 31 January 2026; the \$60 price per thousand impressions was first reported by The Information.
<https://techcrunch.com/2026/02/09/chatgpt-rolls-out-ads/>
-
- [18] **Markey Probes AI Companies on Their Plans to Roll Out Advertising in AI Chatbots**
Office of Senator Ed Markey · 22 January 2026
The letter to OpenAI, Anthropic, Google, Meta, Microsoft, Snap and xAI.
<https://www.markey.senate.gov/news/press-releases/markey-probes-ai-companies-on-their-plans-to-roll-out-advertising-in-ai-chatbots>
-

[19] Many teens are turning to AI chatbots for friendship and emotional support

American Psychological Association · Monitor on Psychology · October 2025
The research the letter leans on: teenagers using chatbots for friendship and support.
<https://www.apa.org/monitor/2025/10/technology-youth-friendships>

[20] OpenAI delays ChatGPT's 'adult mode' again

TechCrunch · 7 March 2026
The feature announced for December 2025, still not shipped in March 2026.
<https://techcrunch.com/2026/03/07/openai-delays-chatgpts-adult-mode-again/>

[21] California Superior Court Consolidates Product Liability Actions Against OpenAI

Hunton Andrews Kurth · 20 May 2026
The coordinated proceeding of February 2026, In re: ChatGPT Product Liability Cases (JCCP No. 5431).
<https://www.hunton.com/hunton-retail-law-resource/california-superior-court-consolidates-product-liability-actions-against-openai>

ON THE EVIDENCE · CORRECTIONS

Not a claim that a court has found OpenAI responsible for the death of the boy from California — no court has found anything, and the case is open. Not a claim that the people who stayed were cowards, and not a claim that the clauses were written to hide a specific harm. 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 4

The Hidden Threat Inside OpenAI: Why They Don't Want You To Know! — 19:56, live since 30 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/tKg52WvHQFE →

facelessyoutubeautomationhq.com/resources/the-other-side-04-the-silence/ →