EXHIBIT 132



HIGHLY CONFIDENTIAL
ATTORNEYS' EYES ONLY

Transcript of Vinay Rao, Ph.D., Designated Representative

Date: November 7, 2025
Case: Concord Music Group -v- Anthropic

Planet Depos
Phone: 888.433.3767
Email: transcripts@planetdepos.com
www.planetdepos.com

1

IN THE UNITED STATES DISTRICT COURT
NORTHERN DISTRICT OF CALIFORNIA
SAN JOSE DIVISION

CONCORD MUSIC GROUP, INC., ET AL.,)
Plaintiffs,) Case Number
v.) 5:24-cv-03811-EKL
) SVK
ANTHROPIC PBC,)
Defendant.)

HIGHLY CONFIDENTIAL--ATTORNEYS' EYES ONLY

VIDEOCONFERENCED and VIDEO-RECORDED
DEPOSITION of
Anthropic PBC, by and through its Designated
Representative,
VINAY RAO, Ph.D.
San Francisco, California 94111
Friday, November 7, 2025

Reported Stenographically by:
MARY J. GOFF
CSR No. 13427
WA CSR No. 21030779
Job No. 606055A
PAGES 1-159

2

HIGHLY CONFIDENTIAL--ATTORNEYS' EYES ONLY

VIDEOCONFERENCED and VIDEO-RECORDED DEPOSITION of VINAY RAO, Ph.D., Volume I, taken on behalf of Plaintiff Publishers, Oppenheim + Zebrak, LLP, at WilmerHale, 50 California Street, Suite 3600, San Francisco, California 94111, beginning at 9:18 a.m. and ending at 2:54 p.m., on Friday, November 7, 2025, before MARY J. GOFF, California Certified Shorthand Reporter No. 13427 and WA CSR No. 21030779.

3

APPEARANCES:

For Plaintiff Publishers
Oppenheim + Zebrak, LLP
BY: COREY MILLER
TIMOTHY S. CHUNG (NYC office)
Attorneys at Law
4530 Wisconsin Avenue
5th Floor
Washington, D.C. 20016
202-957-5043
corey@oandzlaw.com
tchung@oandzlaw.com

For Defendant Anthropic PBC
WilmerHale
BY: TAYLOR GOOCH
Attorney at Law
50 California Street
Suite 3600
San Francisco, California 94111
taylor.gooch@wilmerhale.com
628-235-1026

ALSO PRESENT: Ian Chen, In-house counsel Anthropic
Ellen Stephano, Esquire, WilmerHale (remote Boston)
Videographer: Philip Astor

4

INDEX

WITNESS EXAMINATION
VINAY RAO, Ph.D.
Volume I

ATTORNEY MILLER 7, 161
ATTORNEY GOOCH 149

NUMBER DESCRIPTION PAGE
Exhibit 80 Usage Policy Anthropic_000437165-170 Tab 47 21
Exhibit 81 Anthropic's Consumer Terms of Service, 11-28-2025 Anthropic_0000437176-190 Tab 48 24
Exhibit 82 Acceptable Use Policy Anthropic_0000013675-679 Tab 46 27
Exhibit 83 4-27-23 email, Sanderford Tab 4 Anthropic_00000010907 37
Exhibit 84 CLAUDE's Constitution 5-9-23 Anthropic_000000485-498 Tab 13 54
Exhibit 85 Defendant Anthropic PBC's Supplemental Responses and Objections to Plaintiffs' First Set of Interrogatories Number 12 Tab 78 59

Pages 17-20

17

is required. What I do know the people have to click "Accept Service" requests in -- when they sign up with CLAUDE.AI.

But Anthropic's services are available across API, AWS, and -- Amazon Web Services and Google Cloud. I have -- I have not signed up through those services, so I don't know what exactly that looks like.

Q Does Anthropic expect that all users of CLAUDE will abide by Claude's Terms of Use?

ATTORNEY GOOCH: Objection, scope.

A The question around Terms of Use is -- was not exactly the purview of my team, so there could be a person that is better placed to answer that.

Q (BY ATTORNEY MILLER) Are you familiar with the term "Acceptable Use Policy"?

A Yes.

Q Does Anthropic have an Acceptable Use Policy?

A Yes.

Q Is the Acceptable Use Policy part of Anthropic's Terms of Use?

A Yes.

Q Why does Anthropic have an Acceptable Use Policy?

18

A The -- the models can do a variety of things in terms of asking prompts, and they will respond.

And Anthropic's expectations -- expectation of the users is they work within certain boundaries and not use the model capabilities outside of that or access the model from outside of that, so that is the reason why they exist.

Q So is it fair to say then that the model is capable of doing certain things and Anthropic doesn't want its users to use the model to do certain things that the model is capable of?

ATTORNEY GOOCH: Objection, vague.

A So let me clarify. The expectation is the users follow what is said in the Acceptable Use Policy to be able to continue using the model.

Q (BY ATTORNEY MILLER) What happens if a user does not abide by the Acceptable Use Policy?

A The Anthropic teams and my team in particular has a set of detections which are set up to detect if there is usage that is likely to fall outside the Acceptable Use Policy.

And if that is detected, then a set of interventions are available depending on the type of exception that was flagged, and so those

19

interventions would be applied.

Q What kind of interventions -- withdrawn.

What are the different types of interventions that are available to Anthropic when Anthropic detects a usage of CLAUDE that is outside the Acceptable Use Policy?

A So to be able to answer that, I have to sort of take a step back and explain what the detections are, because certain detections -- certain inventions apply to only certain detections.

So we have mechanisms to detect if a prompt is likely to elicit -- or there is -- elicit something that is likely to be outside the Acceptable Use Policy.

And because this has been detected at a prompt level, there are inventions like prompt modification, which is an -- an extra statement is added to the prompt that provides additional instruction to the model to make sure that the completion is within the Acceptable Use Policy, so in that way we prevent the model from streaming a harmful response.

There are completion-level detections which would be looking at the completion as it is streaming.

20

And in the case of violations that are potentially -- what -- if there are -- if there is content that is potentially violating the Acceptable Use Policy, then we will terminate the response, so stop the response and send a message back to the user.

And then on top of this there are ways in which the teams will then aggregate the behavior of a user, and that could result in user-level actions where you could prevent -- ban the user. That's a term that they use, "ban the user" from using CLAUDE in general.

There is something very specific that's done for child safety. If you want me to go into it, I can go into that as well.

Q We can set the child safety aside for the moment. Under what circumstances will CLAUDE -- withdrawn.

Under what circumstances will Anthropic ban a user for violations of the Acceptable Use Policy?

A So I -- I will have to get you all of the precise reasons, but broadly if there is a belief that a user is in a more -- in a single-minded way trying to circumvent the Acceptable Use Policy, then

Pages 69-72

69

Q (BY ATTORNEY MILLER) But right sitting here right now, you don't know whether that data exists; is that correct?
ATTORNEY GOOCH: Same objection.
A So again, when we talk about data existing in kind of any system, there are mechanisms that you do. You run a test, and that test has multiple objectives; and it -- you can draw multiple inferences depending on how the test is done. And also, there are a bunch of, you know, conclusions that are developed and written down.
What exactly was done to -- whether it answers your specific question is like two different people. Like, this is a question, and some other data may have been done -- computed and so on. So it is probably best to go back to the primary sources and look at it, if there is data that supports the precise question that you're asking.
Q (BY ATTORNEY MILLER) Do you know whether Anthropic has done any tests or studies to determine how effective constitutional -- excuse me -- to determine how effective fine-tuning on a constitutional copyright principle is at preventing users from writing creative prompts to get around more traditional auxiliary systems?

70

A I'll say the same thing: To get a precise answer to your question, we will have to go back to the -- to whatever evaluations that were done and if they support the statement that you are making.
Q Okay. Do you see at the end of line 20, this response cites to an "artifact"?
A Is this the statement -- the line that says "constitutional.r1 Artifact" --
Q I'll ask you not to read the actual citation to keep it out of the record.
A Oh, I see. Okay. Sorry.
Q But do you see that citation?
A Yes.
Q And is that a citation to a portion of CLAUDE source code?
A Yes, so I spoke with Kevin Kelley about this. So from my understanding of what he said is that this is a subset of the Git repo where this update is recorded.
THE COURT STENOGRAPHER: The Git? Subset of the?
A Oh. GitHub repository where this change was recorded.
Q (BY ATTORNEY MILLER) And who is Kevin Kelley?

71

A Kevin Kelley is a member of technical staff at Anthropic.
Q Do you know what his specific role is at Anthropic?
ATTORNEY GOOCH: Objection, calls for speculation.
A It's a specific question. I think it's best to check.
He was the person that I spoke to about the times and dates of the -- when these changes were being made as part of preparation of -- for various defenses that Anthropic had put up.
Q (BY ATTORNEY MILLER) Have you, yourself, reviewed the -- the source code cited in line 20 on page 10?
A This particular source code, no.
Q Prior to your preparation for this deposition in your conversation with Mr. Kelley, were you aware that this particular artifact cited on line 20 existed?
A So I'll go back to what I had answered before we got into this document.
I had -- I have -- I had read the constitution as part of my interview preparation, so I had a broad understanding. Do a line-by-line and

72

word-by-word precision, I don't remember.
Q Are you familiar with something at Anthropic called the "content compare system"?
A Yes.
Q Are you familiar with something at Anthropic called CopyFairy?
A Yes.
Q What is the relationship between the content compare system and CopyFairy?
A So they are -- the outcome of what it sought to do, they are the same.
The idea is to look at completions as they're being streamed and detect if there is a regurgitation of copyrighted material so that way they are similar in terms of how they -- what they do.
In terms of how they are -- how the engineering implementation is, the cease cont -- cont -- content completion system is an updated version of CopyFairy.
Q Why did Anthropic develop the content compare system?
A Just to clarify, we are talking about -- specifically about CCS not CopyFairy.
Q Thanks for that clarification.

Pages 129-136

129

on the first page of Anthropic --
A Hang on. Let me read this.
Is this just two pages?
Q It should be three pages.
A Okay. Three pages. Got it.
Q And for the record, the Bates Number on the first page is Anthropic_0000350023.
And I'm not going to ask you about the whole thing, just about the one part at the beginning.
A Okay. Give me a second. It keeps scrolling down. Sorry. Okay.
Q So what's been marked as Plaintiffs' Exhibit 89 appear to be the record of a Slack chat from February 13, 2024?
ATTORNEY GOOCH: Objection, calls for speculation.
A Okay. I couldn't tell.
Q (BY ATTORNEY MILLER) Do you see the second -- on page 1, the second entry appears to have been written by someone name, "samm."
A I see that.
Q Do you know who that is?
A Yes.
Q Who is that?

130

A An engineer on my team, Sam Martin.
Q And that message appears to have been posted or sent on February 13, 2024?
A Yeah, that's the time stamp next to it.
Q And Mr. Martin writes:
[redacted]
Do you see that?
A Yes.
Q "1P API" means first-party API?
A Yes.
Q And the next line states:
[redacted]
Do you see that?
A Yes.
[redacted]
Do you see that?
A Yes.

131

[redacted]
ATTORNEY GOOCH: Objection, calls for speculation.
A Yeah, for -- to get exactly how it was computed for this, we'll have to check.
From an engineering perspective, this is a fairly standard way of making an estimation of the business impact of a launch.
[redacted]
ATTORNEY MILLER: Tab 38.
(Exhibit 90 was marked for identification and is attached to the transcript.)
Q (BY ATTORNEY MILLER) All right. In the

132

Zoom you should see what I have marked as Exhibit 90, which has Bates Number Anthropic_0000352753.
Do you have this exhibit in front of you?
A Yes.
Q And do you see at the top there's a message from "samm," correct?
A Yes.
Q And that's Sam Martin again, right?
A Yes.
Q And this one is dated February 15, 2024, correct?
A Yes.
Q And Mr. Martin writes:
[redacted]
Do you see that?
A Yes.
Q And then he writes:
[redacted]

133

[redacted]
Do you see that?
A Yes.
Q So when Anthropic enabled Prompt Shield on
[redacted] of Prompt
Shield?
A Yes.
Q [redacted]
A So the -- the system that has been put in place is -- is a fairly hard to -- like when I say "system," it's a business decision system is -- we want the -- both the Prompt Shield and CCS to be turned on for all users all the time.
[redacted]

134

[redacted]
ATTORNEY MILLER: All right. Can we go off the record?
ATTORNEY GOOCH: Sure.
THE VIDEOGRAPHER: We are going off the record at 14:03 p.m.
(A break was taken from 2:03 p.m. to 2:14 p.m.)
THE VIDEOGRAPHER: We are back on the record at 14:14 p.m.
ATTORNEY MILLER: What number are we, 91?
Q (BY ATTORNEY MILLER) Dr. Rao --
A Yes.
Q -- do you have in front of you Exhibit -- Plaintiffs' Exhibit 91, which bears the Bates Number Anthropic_0000012694.
(Exhibit 91 was marked for identification and is attached to the transcript.)
A Yes. Give me a second to read it. Okay.
Q Go to the top the document?
A Yeah.

135

Q This document is titled "Prompt Shield Explainer", correct?
A Yes.
Q And it appears to have been drafted by Nathan Bailey on February 19, 2024?
A Yes.
Q What is this document?
A This is a -- as the document says, it's an explainer of how Prompt Shield is built and -- and what it does. And I think it also has a little bit of a -- sort of a strategic view of what it can do.
Q Okay. You can put that aside. Tab 77.
(Exhibit 92 was marked for identification and is attached to the transcript.)
Q (BY ATTORNEY MILLER) You should see in front of you Exhibit 92 -- Plaintiffs' Exhibit 92, which bears the Bates Number Anthropic_00000440820.
Do you have this document in front of you?
A This is the one that says "CCS ChangeLog"?
Q Yes.
A Yes.
Q Are you familiar with this document?
A Yes.
Q What is this document?
A This is a set of actions that have been -- or changes that have been made to the CCS system.
Q Okay. If you can go to page 5, please, which I think is the last page.
A The one that says "Useful for Archaeology"?
Q Yes.
A Okay. Yes.
Q Do you see a reference to the "StatSig CCS Config"?
A Yes.
Q What is that?
A So StatSig is an engineering system, and it works as follows. So let's say you have a bunch of different systems that need certain numbers that you want to retain some level of configurability over.
So as an example, you say: I want to keep -- I want to run this classifier at this temperature.
Instead of hardwiring there, you keep it in StatSig and configure the classifier to read it off of that so that you can change the temperature at which the model operates.
So all our production system will have -- the production system working whatever. The

136