

TRANSCRIPT OF PROCEEDINGS

In the Matter of:)
)
SECTION 1201 PUBLIC HEARING:)
PROPOSED CLASS 4)
COMPUTER PROGRAMS -)
GENERATIVE AI RESEARCH)
)

Pages: 1 through 94
Place: Washington, D.C.
Date: April 17, 2024

HERITAGE REPORTING CORPORATION

Official Reporters
1220 L Street, N.W., Suite 206
Washington, D.C. 20005-4018
(202) 628-4888
contracts@hrcourtreporters.com

UNITED STATES COPYRIGHT OFFICE

In the Matter of:)
)
SECTION 1201 PUBLIC HEARING:)
PROPOSED CLASS 4)
COMPUTER PROGRAMS -)
GENERATIVE AI RESEARCH)
)

Suite 206
Heritage Reporting Corporation
1220 L Street, NW
Washington, D.C.

Wednesday,
April 17, 2024

The parties convened remotely, pursuant to notice,
at 2:35 p.m.

PARTICIPANTS:

Government Representatives:

EMILY CHAPUIS, U.S. Copyright Office
BRANDY KARL, U.S. Copyright Office
MELINDA KERN, U.S. Copyright Office
KEVIN LI, National Telecommunications
and Information Administration

Panelists:

MICHAEL B. AYERS, AACS LA
ILONA COHEN, HackerOne
AMIT ELAZARI, OpenPolicy
STEVEN R. ENGLUND, Joint Creators and
Copyright Owners
HARLEY GEIGER, Hacking Policy Council
JOSH HARGUESS, Cranium AI
SHAYNE LONGPRE, Massachusetts Institute of
Technology (Ph.D. Student)
MORGAN REED, ACT | The App Association
DAVID JONATHAN TAYLOR, DVD CCA

1 hearings conclude. We ask that everyone speak loudly
2 and clearly and please mute your microphones anytime
3 that you're not speaking.

4 For those of you who are listening in, on
5 Thursday afternoon, we will have a public
6 participation session from 4 to 5 p.m. Anyone who
7 would like to participate in that session can sign up
8 using the link in the chat or on the Copyright Office
9 website. Public comments may relate to any of the
10 classes, but we ask that public participation be
11 limited to three minutes per person.

12 Okay. We'll turn now to Class 4, and let's
13 begin with introductions, starting with the Copyright
14 Office. Melinda, do you want to kick us off?

15 MS. KERN: Hi. My name is Melinda Kern.
16 I'm an Attorney Advisor with the Office of General
17 Counsel.

18 MS. KARL: Hi. This is Brandy Karl. I'm
19 Assistant General Counsel in the Office of General
20 Counsel.

21 MS. CHAPUIS: And we're also joined today by
22 one of our colleagues from NTIA.

23 MR. LI: Good afternoon, everyone. Kevin
24 Li, Special Advisor for AI Policy at NTIA.

25 MS. CHAPUIS: Now I'd like to also invite

1 the participants to introduce themselves, starting
2 with the proponents of the proposed exemption. And
3 when you introduce yourself, will you please state
4 your name and the organization that you're
5 representing? Let's start with Humane Intelligence.

6 (No response.)

7 MS. CHAPUIS: Are they here? HackerOne.

8 MS. COHEN: Hi. I'm Ilona Cohen. I'm the
9 Chief Legal and Policy Officer of HackerOne.

10 MS. CHAPUIS: And OpenPolicy?

11 DR. ELAZARI: Hi, everyone. My name is Amit
12 Elazari. I'm the CEO and co-founder of OpenPolicy.

13 MS. CHAPUIS: Hacking Policy Counsel?

14 MR. GEIGER: Hello. I'm Harley Geiger, and
15 I am the founder and coordinator of the Hacking Policy
16 Council.

17 MS. CHAPUIS: Cranium AI?

18 DR. HARGUESS: Yes. Hello, everyone. Josh
19 Harguess, Chief of AI Security here.

20 MS. CHAPUIS: And MIT?

21 MR. LONGPRE: Hi. I'm Shayne Longpre. I'm
22 a Ph.D. student at MIT conducting research into AI,
23 but I'm here in support of the comments submitted by
24 academic researchers in the field of AI testing and
25 evaluation. Thank you.

1 MS. CHAPUIS: And let's do the opponents of
2 the proposed exemption, please, starting with AACS.

3 MR. AYERS: Hi. Good afternoon, everybody.
4 My name is Michael Ayers. I'm legal counsel for
5 Advanced Access Content System Licensing
6 Administrator, more familiarly known as AACS LA, and
7 we provide content protection technology for Blu-Ray
8 discs.

9 MS. CHAPUIS: Thanks.
10 Joint Creators.

11 MR. ENGLUND: Hi. This is Steve Englund of
12 Jenner & Block, and I'm here representing the
13 Entertainment Software Association, the Motion Picture
14 Association, the News Media Alliance, and the
15 Recording Industry Association of America.

16 MS. CHAPUIS: Okay. And ACT.

17 MR. REED: Hi. My name is Morgan Reed. I
18 am the President of ACT, The App Association.

19 MS. CHAPUIS: And DVD CCA?

20 MR. TAYLOR: Hi. David Taylor, counsel to
21 DVD CCA, which provides licensing technology for CSS,
22 which protect content on DVDs and DVD players.

23 MS. CHAPUIS: Did I miss anyone?

24 (No response.)

25 MS. CHAPUIS: Okay. Great. Thank you all

1 for being here. And with that, I will turn it over to
2 my colleague, Melinda Kern, to start off the questions
3 for Class 4.

4 (No response.)

5 MS. KARL: Okay. I will start off. Could
6 the proponents please provide some examples of
7 scenarios they're trying to address with the proposed
8 exemption? In providing your example, can you please
9 keep the following in mind? What are the copyrighted
10 works that you have in mind for this class? Is it the
11 system, prompt, or something else? What are the types
12 of TPMs that you're concerned about? What are the
13 different circumvention methods you have in mind? How
14 do the activities you have in mind qualify as
15 circumvention? Do these examples also apply to
16 non-generative AI models? If so, how are they
17 different?

18 This question is for supporters.

19 MS. CHAPUIS: I know there's a lot to unpack
20 there, but feel free to take it piece by piece.

21 MS. KARL: Yeah.

22 MR. GEIGER: So I'm happy to speak, but do
23 any of the more operational colleagues want to chime
24 in on that because we're specifically being asked
25 about types of research?

1 MS. KARL: Why don't you get us started on
2 that one? Thank you.

3 MR. GEIGER: Sure. So some of the protected
4 works that we are looking to access through this
5 exemption include broadly computer programs, right,
6 which are a subcategory of literary work and, within
7 that subcategory specifically, the user interface, the
8 code that drives the algorithm, and APIs.

9 The TPMs are the ones that we had cited in
10 our comments. They include account requirements, rate
11 limits, and algorithmic safeguards or so-called
12 guardrails.

13 The particular set of users that we're
14 describing here are persons that are performing good
15 faith research as defined. And so the particular
16 class of works and the specific set of users are
17 similar parameters to what we see in existing
18 exemptions under Section 1201, such as the security
19 testing exemption.

20 So a potential scenario, and, again, I'll
21 leave it to some of my more hands-on keyboard
22 colleagues who perform this research to describe them,
23 but a potential scenario is a researcher that is
24 performing research on discrimination in an AI system.
25 They need an account in order to access that user

1 interface, as well as the code that drives the
2 algorithm, and they engage in prompt engineering,
3 prompt injections, and they lose their account as a
4 result of this. So they become suspended once the AI
5 system operator discovers that they are performing
6 this research.

7 To circumvent their account suspension,
8 which has blocked them from getting access to the
9 protected works, they create a new account. The terms
10 of service forbid this because the terms of service
11 say only one account per user.

12 When they are creating their new account,
13 the circumvention includes the creation of a new
14 username and a password. They may need to use a new
15 email address because their original email address was
16 banned. If there is a subscription, they may need to
17 use a new credit card as well. They may need to use a
18 new IP address, so they use an IP address rotator.
19 But they have circumvented this and created a new
20 account and they're able to continue with their
21 research. So those are -- that is one possible
22 scenario.

23 You asked a question regarding whether or
24 not it is different, whether it is generative AI or
25 not generative AI, and I think the answer in general

1 is no. These TPMs are present in many systems. The
2 protected works at issue are also present in many
3 systems, and the types of research into AI
4 trustworthiness can also apply to non-generative
5 systems.

6 MS. CHAPUIS: Amit, I see you have your hand
7 raised.

8 DR. ELAZARI. Yes. Amit Elazari with
9 OpenPolicy. I'm happy to expand on these comments and
10 agree and support everything Harley just mentioned.

11 So just to kind of provide context on the
12 type of AI auditing mechanisms we have seen, and these
13 have been, you know, broadly documented, including by
14 policymakers in prior work as type of testing methods
15 that are important in order to uncover unintended
16 consequences of AI.

17 So we are familiar, for example, with audit
18 methods that include things like sock puppeting,
19 creation of users in order to exhibit different type
20 of features or type of attributes in order to kind of
21 test the system for potential AI bias in audit. These
22 are well documented. For example, the Sandvig
23 decision that was in the D.C. District in the context
24 of the Computer Fraud and Abuse Act documented some of
25 these auditing methods, and we, in fact, have seen how

1 terms of use can prohibit those type of system and, as
2 Harley mentioned, in combination together with the
3 terms of service and the ability to suspend the
4 account. So the exercise of a technical measure can
5 prevent such type of very useful and important
6 testing.

7 Also, important to note that exactly like in
8 the Executive Order on AI and as policymakers are
9 recognizing, there are a broad set of unintended
10 consequences and there are a broad set of types of AI
11 systems. And AI systems are defined broadly. They're
12 not just generative AI type of systems, but we're
13 really seeing a very, you know, broad definition of AI
14 in policy and, therefore, it's important that the
15 exemption, as we said in the comments, will apply
16 broadly as well.

17 So I think, you know, I am looking to our
18 technical colleagues here on the line to talk a little
19 bit more about their type of research, but we are
20 seeing this intersection between security research and
21 broader safety research and bias research, and there
22 is a broad set of testing that is being done that can
23 be characterized as broader than just traditional
24 security techniques that are needed in order to
25 evaluate the type of unintended consequences of AI

1 that we see today and that would emerge in the future.

2 MS. KERN: Thank you.

3 Mr. Longpre?

4 MR. LONGPRE: Yeah. I'm happy to expand on
5 that a little bit. So I'm not a lawyer, but I am an
6 AI researcher. And in addition to what Harley and
7 Amit said, maybe I can point first to the open letter
8 that was signed by 350 researchers in the field that
9 we cite in our comment, and that letter sort of had
10 three points that seemed to gain broad traction in the
11 community.

12 The first is that this type of research into
13 AI trustworthiness that includes bias, discrimination,
14 misinformation generation, and some other things is
15 really timely and critically important and there isn't
16 enough of it.

17 And the second point is that this good faith
18 research and many of the researchers that are even
19 doing this research are feeling a form of chilling
20 effects because of fear of potential liability for
21 violating terms of service and/or trying to circumvent
22 guardrails or creating new accounts after their
23 accounts have been terminated in order to do this good
24 faith research.

25 And so that community in the letter that was

1 widely signed is supporting broader protections for
2 that type of public interest, in our view, beneficial
3 research.

4 I'll also add, I'm going to talk a little
5 bit more about the types of guardrails if that would
6 be beneficial, but you asked at the end about
7 generative AI versus other types of AI. I'll add
8 that, in our view, this distinction is a little bit
9 artificial. There are many similar systems and models
10 used, for example, for facial recognition that is not
11 a generative model, but it still has very important
12 consequences for society. There's still TPMs. It's
13 still important to evaluate these systems for bias,
14 which there have already been many cases discussed for
15 that particular application. So we think this
16 research is important in both those places.

17 MS. KERN: Thank you.

18 We'll go Mr. Taylor, then Mr. Harguess,
19 please.

20 MR. TAYLOR: Yes. Thank you. I think it's
21 very important to ascertain what is the circumvention
22 that's going on here. And what I've heard that is
23 traditionally understood to be 1201 access control is
24 only the use of password and the proverbial walled
25 garden. And in terms of a 1201 act of circumvention,

1 the only way that I'm familiar with an act of
2 circumvention being really in the terms of a password
3 is a brute force attack. And when I read the initial
4 comments and even when I read the reply, I did not get
5 the notion that they were going to use brute force
6 attacks for the purposes of gaining access to whatever
7 they may mean by generative AI.

8 And so terms of use that they may violate,
9 those aren't governed by 1201 and this rulemaking
10 really has no ability to address that.

11 MS. KERN: Thank you.

12 And I apologize. Dr. Harguess.

13 DR. HARGUESS: All good. So, yeah. I think,
14 in the original example, I do want to support, you
15 know, kind of everything that was said there. That
16 scenario of being, you know, kicked out of an account
17 while you're doing prompt injection, some of this is
18 viewed from a lens of red teaming. I know that was
19 submitted also as a concept that we'd like to be, you
20 know, as part of this. Red teaming can uncover, you
21 know, things within security, but it can also uncover
22 things in trustworthiness, bias, you know, other types
23 of things that come out of these models.

24 Agree with everyone that there is no
25 distinction between AI and generative AI. You know, a

1 year from now, two years from now, we may not be
2 having this generative AI discussion. It may be a
3 very different discussion, so I do want to make sure
4 that those lines are clear.

5 And further, you know, this idea of kind of
6 red teaming, all models - all AI models are
7 susceptible to some type of attack. They can be
8 broken and manipulated. This ability to be able to go
9 in and do, you know, sort of this red teaming or, you
10 know, this analysis, this research onto these models
11 so that we can better understand the landscape from a
12 security perspective, from all of these other
13 assurance perspectives, is really important. It
14 informs the community. There's things like MITRE
15 ATLAS, which collects, you know, a lot of these
16 security incidents and these different tactics and
17 procedures.

18 There's things like OWASP. They're trying
19 to understand, you know, what are the top 10, you
20 know, items that you need to care about when we're
21 thinking about generative AI and other types of
22 machine learning.

23 So these types of activities are really
24 important and so we just want to make sure that, you
25 know, researchers and practitioners that are trying to

1 inform the community about these types of AI assurance
2 issues are able to do those jobs.

3 MS. KERN: Thank you.

4 Mr. Reed?

5 MR. REED: Hi. Thank you. I want to try to
6 clarify one thing, which I think -- and it's funny, I
7 was looking at Harley's window. I think I can almost
8 see my building through his window over there. I'll
9 wave at him, he's across the street, as a former
10 Venable person.

11 Here's the thing that I thought needs to be
12 clarified. I don't believe that Harley's asking for
13 this, but I don't think the proponents of this -- are
14 you arguing that companies should not be able to block
15 an unknown hacker? Because, if you don't contact the
16 company in advance to tell them that you're red
17 teaming, you are essentially a potentially malicious
18 hacker.

19 And while you're performing good faith
20 research, I would say that it's good security
21 practices to block a person from using the same credit
22 card, to block a person from using the same email.

23 As a former person on the other side of the
24 table, I would be remiss in my duties if I wasn't
25 implementing every possible barrier. So I want to

1 clarify it's -- that the proponents are saying they
2 don't want to face a copyright consequence from taking
3 this action, or do you think that there should not be
4 TPMs preventing you from doing those activities? And
5 I just want to clarify that because that's a very
6 different take than the idea of, well, we shouldn't
7 face a copyright consequence or a lawsuit after the
8 fact.

9 MS. KERN: Thank you.

10 Mr. Englund, please.

11 MR. ENGLUND: So, after hearing the
12 proponents identify a number of scenarios in response
13 to the Office's original question, I think it's
14 important to observe that the proposal that has been
15 put forth in regulatory language by HPC and all the
16 comments, including some of the ones just in the last
17 few minutes, asking for an exemption that is wildly
18 broader than the scenarios that have been just
19 identified, once we move beyond generative AI, AI is
20 ubiquitous. And so, as it affects my clients, we have
21 comments from DVD CCA and AACCS talking about hacking
22 the software on DVD players and Blu-Ray players, but I
23 think we're talking about breaking the TPMs on video
24 games that have AI features. I think we are talking
25 about circumventing the TPMs that provide user

1 authentication for streaming services like Spotify and
2 Netflix and new sites in social media that have
3 recommendation engines powered by AI.

4 And it's not hard to think of lots of other
5 systems out in the world that are powered by AI:
6 credit card fraud prevention, autonomous vehicles, you
7 name it. And so we should all be clear that when we
8 say we want to be able to conduct testing on all AI,
9 we're talking about a tremendous range of things.

10 And in terms of the scenarios themselves, I
11 think somebody suggested it's obviously not in the
12 Office's power to immunize users from terms of use
13 violations or prevent account suspensions. And it's
14 not entirely clear as a general matter whether
15 everything that's been talked about here is a
16 circumvention, but important to recognize that there
17 are very good reasons for online services to implement
18 account authentication to ensure that users of
19 subscription services, for example, are who they say
20 they are and they're using the services in the way
21 that they paid for.

22 Beyond that, the rate limitations and
23 limitations on multiple accounts serve important
24 purposes of allocating system usage. And so, if
25 people are exceeding those limitations, they may be a

1 malicious actor that may be taking scarce resources
2 away from other users, and that's a problem that
3 shouldn't be ignored here.

4 So this is a very different proposal from
5 the kinds of security research proposals the Office
6 has considered in the past that don't have
7 implications for online services or other users of
8 online services and seem particularly inappropriate
9 where AI is incidental to a service, particularly one
10 providing access to creative content.

11 MS. KERN: Thank you.

12 Mr. Geiger, please.

13 MR. GEIGER: Thank you. I'd like to respond
14 to the three opponents.

15 So, first, on the question of whether or not
16 the proposal is forbidding service providers from
17 blocking "unknown hackers," I would argue that this is
18 a very serious misunderstanding of the law and a
19 misreading of the plain language of Section 1201 and
20 our exemption.

21 As noted, Section 1201 does not prohibit a
22 service provider or the owner/operator of a computer
23 program or owner of a protected work from taking steps
24 like suspending accounts. That is not the issue.
25 It's not what our exemption or really any exemption

1 proposed under Section 1201 would do.

2 Second, on brute forcing passwords, this too
3 I think is a serious misunderstanding of the law. And
4 I think that we should avoid hyperbolic diversions.
5 Section 1201(a)(3) notes that circumventing a
6 technological measure is bypassing or avoiding a
7 technological measure. So passwords in the scenarios
8 that we described are involved, but as described, the
9 bypassing or the avoiding of that technological
10 measure does not have to involve brute forcing
11 passwords.

12 Lastly, Mr. Englund described a range of
13 software that may be covered by our proposed
14 exemption. He's correct in all of that. What we are
15 proposing is an exemption that is cabined to computer
16 programs that run IA systems. This is actually
17 narrower than several existing exemptions.

18 Most existing exemptions apply to computer
19 programs. So Section 1201 has an exemption for
20 security testing, for encryption, for reverse
21 engineering, and they apply to computer programs.
22 Encryption runs on more things than AI. Security
23 issues are inherent in all software, not just computer
24 programs that run artificial intelligence.

25 So, again, I think that the boundaries of

1 our proposed exemption are not overbroad and, in fact,
2 are narrower than several existing exemptions. Thank
3 you.

4 MS. KERN: Thank you.

5 All right, Mr. Ayers?

6 MR. AYERS: Thank you. Actually, responding
7 also to Mr. Geiger but piggybacking a bit on Mr.
8 Englund. AACS LA actually is very concerned for some
9 of the reasons that Mr. Englund mentioned. And, Mr.
10 Geiger, to the extent that you're saying that the
11 proposal is much narrower than we may be perceiving, I
12 think that's helpful, but I would not necessarily
13 point to a misstatement of the law on the parts of the
14 opponents as much as it is perhaps a failure of the
15 proponents to have made a clear proposal about what is
16 needed and what is actually on the table.

17 Our concern does extend to -- especially
18 having clarified today that referring to AI is not
19 exclusive to generative AI, that it includes other AI
20 tools that are not arguably classified as generative.
21 And so, to what extent does that impact a Blu-Ray
22 player in which a manufacturer has incorporated an
23 up-res'ing tool that might be considered an AI
24 application for the purposes of taking a lower
25 resolution piece of audiovisual content and presenting

1 it in a attractive higher resolution form?

2 To what extent is there a concern about how
3 the implication -- or how the application of that AI
4 tool impacts, to the extent there's no racial
5 differences in how people are presented in the
6 up-res'd content? You know, so does that suddenly
7 mean that this Blu-Ray player in question and perhaps
8 even the disc in the tray being played are now subject
9 to the exemption?

10 So those are the concerns that we're coming
11 here with, and so it's very helpful to hear that the
12 proponents would like us to perceive their proposal as
13 narrow. I don't think it's as narrow as you think it
14 is as currently proposed. So I would propose that we
15 look at actually specifying a little more so we can
16 clarify what's actually on the table.

17 MS. KERN: Thank you. And, unfortunately,
18 I'm going to have to invoke my moderator discretion
19 here. This question was very general and we're glad
20 everyone got a chance to answer it. But I would like
21 to move on and pass the mic over to Kevin Li.

22 So, Kevin?

23 MR. LI: Thank you very much. I'd like to
24 dig deeper into the question of what specific TPMs and
25 underlying protected works could be at issue in this

1 exemption, you know, and particularly, and this is a
2 question for both proponents and opponents, but I'd
3 like to start with the proponents. I'd like to pose a
4 couple hypotheticals. In particular, I'm, you know,
5 mostly drawing off of the three categories of TPMs
6 identified by the Hacking Policy Council in their
7 reply comments.

8 And, first, I'd like to start with to what
9 extent -- you know, what is included in the computer
10 programs. For example, you know, if an AI system is
11 trained on a system prompt or an instruction prompt
12 and, you know, someone uses prompt injection, a
13 researcher uses prompt injection to try to obtain that
14 copyrighted system prompt, is that considered part of
15 the computer program at issue?

16 If the model weights themselves have
17 memorized some underlying copyrighted work and
18 regurgitates that copyrighted work in response to some
19 kind of system prompt or some kind of adversarial
20 prompt, is that part of the computer program that is
21 being circumvented?

22 And, secondarily, I'd like for you to
23 discuss in more detail what the algorithmic safeguards
24 at question could be.

25 MR. LONGPRE: Maybe I can start. Kevin,

1 thanks for the question. I think that -- I'm not a
2 copyright expert. I'm a researcher again, but I'll
3 defer some of those questions to others. But, when
4 interacting with these programs and doing research on
5 them, we are probing them in various ways to
6 investigate them through the interface, like the
7 playground that some of these models have. We're also
8 investigating them through an API and probing
9 different parts of the system, including the various
10 filters and moderation on the inputs and also on the
11 outputs. As I understand it, those algorithms and
12 elements of the system and software that govern that
13 may be the copyright material that we're interacting
14 with and investigating.

15 If I can address something really quickly
16 about a prior comment that was made about how some of
17 the research might not be overloading the system or
18 not paying or something whereas real customers do, all
19 the research that I've seen and the people I've been
20 speaking to are having their accounts suspended and
21 have fear of liability when they're paying for their
22 accounts and they are using the regular systems as
23 they're meant to be used, except they're doing good
24 faith academic research. And so that's not, you know,
25 shirking costs the company has or something. And I

1 think there's an important distinction there because,
2 otherwise, it wouldn't be good faith.

3 MR. LI: And perhaps we could go now to Mr.
4 Geiger.

5 MR. GEIGER: Sure. Mr. Li, do I understand
6 your question correctly? Are you asking if the output
7 is one of the protected works that we are seeking
8 access to?

9 MR. LI: I think that, you know, I would
10 like to get clarification on what particular elements
11 of -- you know, whether, for example, you now,
12 information that is contained within the AI model
13 itself, you know, do you view that as part of the
14 protected work, the computer program.

15 MR. GEIGER: Like the training data?

16 MR. LI: Or information that the model has
17 learned from the training data.

18 MR. GEIGER: So, largely, I think that
19 question is moot, honestly. I think that we are
20 seeking to - I understand that the output of the AI
21 system may or may not be copyrighted. In some cases,
22 the AI system owner explicitly says that they are
23 relinquishing copyrights to some of the output,
24 particularly for generative models.

25 The training data, likewise, you know, that

1 can be a complex land as to whether or not that is
2 protected. I think, for purposes of the research, the
3 protected works that we are talking about accessing
4 are the user interface. So, there, you know, you log
5 onto an AI system to engage with it. It is the
6 software that you're viewing once you have gotten past
7 your login window, which is the technological access
8 barrier, and then, number two, the software that
9 drives your engagement with the algorithm and the
10 software that drives the algorithm itself, so the code
11 that is enabling the algorithm to work, which, again,
12 as a computer program is a subcategory of literary
13 works.

14 And then, lastly, there are forms of
15 research that are undertaken on the APIs of artificial
16 intelligence systems as they appear in other
17 instances, so as they are licensed in other places.
18 So those three computer programs or, you know,
19 examples of computer programs are what we are seeking
20 access to, especially with this exemption.

21 I think you had also asked about guardrails.
22 Would you mind rephrasing your question, please?

23 MR. LI: Yeah. In the reply comments, the
24 Hacking Policy Council discusses algorithmic
25 safeguards as one category of potential TPMs that

1 would need to be -- that this exemption proposes to
2 allow circumvention of, and that discussion, I would
3 appreciate going a bit deeper on that.

4 MR. GEIGER: Sure. So this is just one
5 technique for artificial intelligence research, which
6 is circumventing our guardrails, which are algorithmic
7 safeguards that prevent the or are designed to prevent
8 the AI system from engaging in activity like producing
9 harmful content or engaging in bias. And part of the
10 purpose of the research is to essentially circumvent
11 those guardrails and, in doing so, they may access
12 other features or enable the AI system to operate in
13 different ways.

14 So, for example, some research circumvents
15 these guardrails by elevating user privileges, so,
16 essentially, convincing the AI system that you are an
17 administrator and, therefore, the guardrails no longer
18 apply to you as a user, giving you greater access to
19 unfiltered responses from the algorithm.

20 There is one thing I'd like to just
21 highlight from an earlier answer from Mr. Ayers. I
22 thought that you had a brilliant example of artificial
23 intelligence research that I would like to just
24 highlight as falling under our exemption.

25 You had mentioned a Blu-Ray player with an

1 artificial intelligence tool that can up-res content
2 and the AI research would identify potential racial
3 disparities with that up-res'ing, and I think that
4 that is a terrific example. That research would not
5 violate copyright. I think that research would be
6 socially beneficial and is exactly the kind of
7 research that we would envision as being encompassed
8 by our exemption, and that research would then enable
9 further future up-res tool-makers to avoid racial
10 disparities in their tool. So I do appreciate that
11 example and think that that is a great one for
12 purposes of discussion.

13 MR. LI: Thank you, Mr. Geiger. I'm going
14 to give Ms. Cohen a chance to speak as well, and then
15 I'd like to hear from the opponents.

16 MS. COHEN: Thank you so much. I just
17 wanted to mention, so HackerOne is a global leader in
18 human-powered security but also trustworthiness
19 testing. So I just wanted to try to answer your
20 question about how we might try to seek to bypass
21 algometric safeguards.

22 And so, you know, oftentimes, the
23 researchers that we will work with will seek to assess
24 the behavior of the model to understand sort of the
25 rare instances in which we can get a system to display

1 inappropriate content or other undesirable outputs so
2 that the underlying owner of the AI system can
3 ultimately fix those undesirable outputs for future
4 use. That's a circumstance in which we would
5 intentionally try to bypass any algorithmic safeguard.

6 MR. LI: Thank you, Ms. Cohen.

7 Now I'd like to give the opponents a chance
8 to address in particular, if you could address whether
9 the algorithmic safeguards in your view count as
10 technological protective measures under 1201, as well
11 as, you know, whether anything else on the previous
12 comments that you'd like to address. Mr. Reed?

13 MR. REED: Thank you. It's an interesting
14 situation to find myself in one where, in general, in
15 the larger scope, I agree with the proponents in the
16 sense that bias testing is really important. It's
17 actually pretty critical. It's something that we have
18 our own set of policies on how we should do it. It's
19 the question of venue and is 1201 the right vehicle.

20 The thing that I'm struggling with right now
21 is, in Mr. Geiger's recent example, I'm trying to
22 figure out what he wants to Copyright Office to do?
23 Because, as he established and is the law well known,
24 systems have a right to essentially protect themselves
25 through TPMs. And is he envisioning a world in which

1 good faith researchers receive a token from the large
2 language model or the foundation model level and then
3 that makes it clear that other normal security
4 procedures can be appropriately bypassed to continue
5 the research?

6 At this point in time, I'm not sure how that
7 works because, without a token, we're going to throw
8 everything against the wall to make sure they can't
9 break in, and that's the point that I want to go with.

10 The second part that came up earlier and I
11 think is really critical is -- and let's face it.
12 Kevin, we all know this very well. AI is essentially
13 a marketing term because, in a lot of ways, a rules
14 engine, a sufficiently sophisticated rules engine is
15 essentially sometimes classified as AI.

16 So I think, with the expansion of this,
17 moving it from generative AI or predictive AI into a
18 lot of other fields is going to really open up
19 questions that I haven't even considered.

20 When this first happened, this was really
21 around generative AI, and I was hoping we could look
22 for kind of narrow scope opportunities to solve the
23 situation. But, if it opens up to things like rules
24 engines and kind of the whole panoply of general use
25 software, it raises a whole lot of questions.

1 So, to your primary answer, I think that's
2 the big question. It doesn't seem like this is the
3 venue for the Copyright Office. Through the EO,
4 there's a lot of other agencies that are working on
5 these exact issues, and we're going to have to figure
6 out how do you provide a good faith research effort
7 with the tools to do the kind of bias research they
8 want to do but not in a way that compromises the
9 security for untoward actors.

10 MR. LI: Let me pass it to Mr. Englund, but
11 before I do that, let me just say that one question
12 that would be helpful to answer is, is it a 1201
13 violation absent an exemption to use --

14 MR. REED: Thank you.

15 MR. LI: -- what the commenter refers to as
16 a jail break prompt to attempt to reveal information
17 about the underlying AI system?

18 MR. ENGLUND: So I raised my hand to respond
19 to the various descriptions of software in response to
20 Mr. Li's original question. And at the risk of
21 repeating a point that I made earlier, I think it's
22 important to highlight that, once again, the answers
23 that Mr. Li received do not at all resemble the
24 regulatory language that was proposed. I actually
25 have a very difficult time parsing the regulatory

1 language that was proposed.

2 There is a lot of stuff in Proposed
3 Paragraph 1 about what devices circumvention occurs
4 on, and I have a hard time relating that to different
5 kinds of AI systems. But, once you say that it has to
6 occur on the proper devices, the exemption is simply
7 for computer programs solely for the purpose of good
8 faith AI trustworthiness research, any computer under
9 the sun.

10 And so we heard about, well, it's the UI or
11 it's the APIs or it's some other software, but that's
12 not what the exemption says. The exemption says any
13 software under the sun, so, again, it extends to
14 things like the security software on a DVD player. It
15 extends to the user authentication software on a
16 streaming service and anything under the sun that uses
17 AI.

18 MR. LI: With a view to helping us get
19 clarity on what exactly the TPMs are, a question for
20 you, Mr. Englund and then also Mr. Taylor and Mr.
21 Ayers, is, you know, whether the specific example of a
22 jail break prompt would count as a 1201 violation
23 absent an exemption in your view.

24 MR. ENGLUND: I don't think the proponents
25 have made enough of a record for me to opine on that

1 question. I think it's just not clear on the record
2 before us.

3 MR. LI: Mr. Taylor?

4 MR. TAYLOR: Yes. So I'll take a stab. I
5 think your first set of questions, I mean, overall, is
6 very good because it gets to the heart that there
7 aren't enough examples or any examples of what we're
8 really talking about here. And we just didn't wake up
9 overnight with this good faith security research.
10 Proceeding after proceeding, we had example after
11 example.

12 But, to answer your more recent question,
13 no. I mean, I had to think about this a lot. I had
14 to think about when we say "technical measures," what
15 are we talking about? And 1201 defines technical
16 measures that are protected are those that are access
17 controls and copy controls, and what are actually
18 described here I think could approximate maybe a copy
19 control, but I don't know. I know, in my mind, when I
20 read it, that it's not an access control.

21 And so I had to look at to see what are we
22 talking about here, and the only thing I can kind of
23 come up with is the interactiveness of -- I'm going to
24 step on somebody's toes here, and I apologize -- video
25 games, right? And so what we're talking about maybe

1 is referred to as cheats. When you do something to
2 the video game that enhances the play, we're really
3 talking about, you know, the performance of it. We're
4 talking about, you know, what kind of cheats can be
5 put in place to result in the program acting
6 differently, and I don't think that is an access
7 control and I think that you just look at those
8 examples and you draw the analogy.

9 MR. LI: I'm going to interrupt. I've been
10 asked to move this along. If I could give Mr. Ayers,
11 Mr. Geiger, and Mr. Longpre 15 seconds each. I'm so
12 sorry.

13 MR. AYERS: Sure. Real quick. So just I
14 would note that, yes, I'm skeptical that it would be a
15 circumvention action that's covered in this
16 proceeding. I would also note, though, that we have
17 sort of gone over the added other TPMs to the list
18 when we talk about the research project that Mr.
19 Geiger thought might actually be interesting in that
20 it does impact systems that are encrypted and would
21 require something other than just using a fake ID or
22 going beyond the terms of use.

23 MR. LI: Mr. Geiger?

24 MR. GEIGER: Yes. So you asked is it a 1201
25 violation to use a jail break prompt to reveal

1 information about the underlying system. If that
2 information about the underlying system is a protected
3 work and a guardrail is preventing access to that
4 information and bypassing or avoiding that guardrail
5 via a jail break prompt gives you access to that
6 information about the underlying system, then I would
7 argue yes. And if prompt engineering or jail break
8 prompts, guardrails circumvention, is not a
9 circumventing a technological protection measure, that
10 would be an excellent thing for the opponents and the
11 Copyright Office to clarify in writing.

12 MR. LI: Thank you, Mr. Geiger.

13 MR. GEIGER: Real quick, if I may, because
14 there were a lot questions. What is Mr. Geiger asking
15 the copyright to do? They were describing a token for
16 good faith research, et cetera. Without the token,
17 we'll throw everything against the wall to prevent
18 this unauthorized use.

19 All of that is appropriate actually. And
20 the only thing that we are asking the Copyright Office
21 to do is to create an exemption for good faith AI
22 trustworthiness research under Section 1201 that would
23 shield them from liability for good faith research
24 under Section 1201. And this, the Copyright Office,
25 is an extremely appropriate venue for that request.

1 Indeed, it is the only venue for that request.

2 As far as the scope goes with every type of
3 software under the sun, I would just reiterate again
4 that that is exactly the scope of numerous exemptions
5 that exist right now under Section 1201, computer
6 programs.

7 MR. LI: Mr. Geiger, I'm sorry to interrupt.

8 MR. GEIGER: No problem.

9 MR. LI: We have only a limited amount of
10 time. Mr. Longpre, if you could make a very brief
11 point?

12 MR. LONGPRE: Yeah, I can skip mine. I just
13 wanted to echo what Harley was saying in response to
14 Mr. Reed that we're not asking for a special token or
15 infrastructure. It's an exemption. Yeah.

16 MR. LI: Thank you, Mr. Longpre.

17 I'm going to pass this back to Ms. Kern.

18 MS. KERN: Thank you.

19 I just had a quick question that I hope that
20 ACT/The App Association could please expound upon that
21 was within their comments.

22 So The App Association states that granting
23 the exemption mandates to allow open access to
24 otherwise protected software. Could you please
25 elaborate on that point? And for the other

1 participants that are opposing the proposed exemption
2 as well, what works would be affected if granted? And
3 if everybody can please keep their remarks very short.
4 We're almost down to our last hour, and we have a lot
5 to get through. Thank you.

6 MR. REED: Yes. This is Morgan. I think we
7 covered a lot of it in the back-and-forth that we just
8 had with Mr. Li's question. And I think I appreciate
9 the clarity that Harley and others have provided
10 around what they're asking for, but, to support my
11 fellow opposition, I feel like it has moved around a
12 little bit through the original proposal that we filed
13 against and what we're hearing today.

14 And so the nearest that we've gotten clarity
15 around it, it is post-fact liability protection. So a
16 researcher does the action. They break in. The
17 company whose LLM it was that they went after is
18 unhappy in some way or form or another and goes after
19 them for a copyright breach. So what they really
20 want, what they're really asking for is post-fact
21 liability protection, and that helps clarify it.

22 As far as the systems and what that means is
23 it gets back to the same thing I said to Kevin
24 earlier, Mr. Li earlier, which is, if we open this up
25 not just from generative AI and the idea of foundation

1 models but to basically to all computer programs, I
2 think that the size and the scope of this gets really
3 large to handle.

4 So, to answer your question, most of what we
5 said in the earlier ones, I think it's clear now that
6 the concern we have is this is strictly post-fact
7 liability protection that could be good, could be not,
8 but I'm not sure this is the right venue for it.

9 And then, finally, I am concerned about
10 expanding it from AGI to what amounts to all computer
11 programs. Thank you.

12 MS. KERN: Thank you.

13 Mr. Ayers?

14 MR. AYERS: Hi. Thank you. Yeah. As far
15 as the scope of the works that would be covered, I
16 think our concern is that it could arguably be read to
17 cover not only devices and applications that are
18 involved in the playback of copyrighted audiovisual
19 content, such as a Blu-Ray player or a DVD player,
20 but, under certain circumstances, might also be read
21 to extend to the content being played back, whether
22 it's on a disc, an optical disc in the drive or
23 training material used for an AI device. So that's
24 the expanded scope of works that are covered that
25 would be concerning to us.

1 MS. KERN: Thank you.

2 Ms. Elazari?

3 DR. ELAZARI: Yeah. I just want to take
4 this opportunity to again reiterate and agree with
5 some of the comments made by Mr. Geiger and Shayne.
6 It's important to recognize, and we provided it in our
7 comments, that we can draw on the concept of AI
8 systems as it's being proposed in the Executive Order.

9 In fact, the same Executive Order is
10 proposing that red teaming and such testing of AI is
11 not just appropriate but desirable. So, you know,
12 currently, concepts in policy are being evolved to
13 define AI systems. We can draw on these definitions
14 and the proposal we brought forward for this exemption
15 is, again, building on an existing security exemption
16 and, therefore, as Mr. Geiger suggested, has the
17 already appropriate guardrails in place.

18 So I think it's important to emphasize that
19 this kind of terminology of AI, you know, it's a
20 marketing term, this is the venue to actually consider
21 the implication for the research ecosystem that are
22 being, as Shayne alluded, very well documented about
23 the concerns about liability to address this
24 anti-hacking limitation. This is, in fact, the same
25 venue where the security exemption has been considered

1 as well.

2 MS. KERN: Thank you.

3 Mr. Englund, please.

4 MR. ENGLUND: I agree with Mr. Ayers'
5 remarks a moment ago, and just to expand on them a
6 little bit, it does sound based on the discussion over
7 the last half hour that the proponents are talking
8 about circumvention that could potentially expose
9 creative works that are currently protected by TPMs.

10 It certainly seems like that's the case for
11 video game software that incorporates AI features. I
12 think they're saying they'd like to be able to
13 circumvent the TPMs on those games. And it also
14 sounds like they would like to be able to circumvent
15 user authentication on the streaming services, which
16 potentially exposes the creative works available
17 through those services.

18 MS. KERN: Thank you.

19 Mr. Geiger?

20 MR. GEIGER: Yes. So, just to the point
21 about whether this has moved around from the original
22 proposal, that is flatly incorrect. We have provided
23 very specific language, and from our initial comments,
24 which included that language, to this hearing, the
25 language is strikingly consistent. So the proposal

1 has actually not moved around. Our proposal was not
2 limited to generative AI and was very clear about the
3 types of programs that this would operate under.

4 Again, it is not all computer programs.
5 Anyone that takes a look at the language we've
6 proposed will see that. It is computer programs on a
7 lawfully acquired device or machine on which an AI
8 system operates. And we use a definition of AI
9 systems that is presently in use throughout U.S. law,
10 as well as the recent Executive Order, but not limited
11 to that Executive Order. So this is a narrowly
12 defined class of protected works, as well as a
13 specific subset of users consistent with the Copyright
14 Office's existing exemptions.

15 MS. KERN: Thank you.

16 Mr. Taylor?

17 MR. TAYLOR: Yes. Thank you. I just would
18 like to offer the quick perspective that this
19 rulemaking, the security exemption that we have, the
20 security research exemption that's already in place,
21 it didn't happen overnight. It was Mr. Feldman who
22 came here repeatedly with a bunch of proponents,
23 different times developed a very concrete record in
24 which the Office was able to evaluate the claims, and
25 we don't have such a record here. In fact, we

1 explained to them in our opposition that we don't
2 understand what you're talking about. The reply did
3 nothing more as informing it.

4 The example that we've given that we've
5 actually provided to them during this discussion is
6 very hypothetical and it's not a concrete proposal for
7 this rulemaking to actually recommend an exemption.
8 Thank you.

9 MS. KERN: Thank you.

10 Ms. Cohen, please.

11 MS. COHEN: I'd just like to align myself
12 with Dr. Elazari and Mr. Geiger. The comments that
13 are provided by the Hacking Policy Council provide
14 very clear definitions which the opponents seem to
15 ignore.

16 And in terms of the correct venue, I'd just
17 note that the Department of Justice has weighed in
18 here, providing support for the proponents of this
19 exemption and drawing the conclusion that this very
20 much is the correct venue and the correct action.

21 And in addition, we already have the
22 terrific record that the opponents keep mentioning
23 with respect to security research that helps and
24 informs this action, but it doesn't mean that we need
25 to duplicate that again. We already have the record

1 that we have created, and we're therefore using that
2 in addition to this additional information being
3 provided here.

4 MS. KERN: Thank you.

5 Mr. Reed, I see your hand is up, so I will
6 give you 15 seconds, but then we are moving on.

7 MR. REED: Yep. The Executive Order said
8 that the U.S. Copyright Office should report to the
9 Administration on AI and copyright, but you all have
10 not even released the first of your three reports. So
11 we may take a completely different perspective because
12 you are all currently undergoing your own process
13 around these questions, and I'd love to see what you
14 come up with. Thank you.

15 MS. KERN: Thank you. And I'll pass the mic
16 over to my co-worker, Brandy Karl, please.

17 MS. KARL: Yes. Hi. This one's for
18 OpenPolicy. In the 8th Triennial, with regard to the
19 current exemption on security research, the Office
20 looked at a proponent's request to remove what is
21 known as the access limitation that circumvention be
22 undertaken solely for the purpose of good faith
23 security research.

24 We ultimately concluded that based on our
25 rulemaking record that the access limitation did not

1 create a reasonable risk of chilling good faith
2 security research, and absent specific evidence that
3 the access limitation is likely to chill otherwise
4 protected security research, the Register could not
5 conclude that the language is likely to cause an
6 adverse effect.

7 This cycle, you requested that the Office
8 not include similar language in this exemption's
9 regulatory language if granted. Your comment said,
10 "Since there are concerns regarding the ambiguity
11 associated with the use of the term 'solely' in the
12 security research exemption that this should warrant
13 not including this language."

14 Could you provide the Office with
15 information about how the language would likely chill
16 any security research associated with generative AI
17 research?

18 DR. ELAZARI: Yeah. I'll speak about this
19 briefly, but I would also like to invite, you know,
20 other proponents, so the Hacking Policy Council and
21 HackerOne, that have a lot of experience with this. I
22 would appreciate their response to this.

23 I'm just going to give a concrete example.
24 Today there are concepts of data abuse by bounties and
25 bug bounties for auditing. I myself, in fact, got a

1 bounty like this where a researcher might be doing
2 good faith research, producing some, you know,
3 valuable insight but also getting compensated right
4 after the fact for such activity. So this would be
5 not just solely for good faith research, but there
6 could be a monetary, for example, value.

7 We also know there is a lot of pentesting
8 companies, and I think there is ambiguity about the
9 term "solely," which means potentially, you know, what
10 happens if you're actually conducting testing for
11 pentesting purposes of getting some kind of monetary
12 value, the activity is also producing some value of
13 research and this can be done also in the context of
14 academic research, right, where there is a fellowship
15 or some kind of grant.

16 And so the context of the testing could also
17 involve some kind of other value. So I think, because
18 of that ambiguity that was actually also documented in
19 the context of the security research discussions, at
20 least in some of the comments that I've seen from CDT
21 and others, we propose that there could be a
22 consideration for the removal of the term "solely,"
23 but it's important to know that any -- you know, as
24 the DOJ suggested in their commentary, even without
25 the removal of the term "solely," I think there is a

1 lot of value in creating this exemption.

2 So I would encourage the Copyright Office
3 to, you know, consider the proposals as they are in
4 support of it and yes, consider the idea of removing
5 the term "solely," but moving the conversation along,
6 even if eventually the Copyright Office decides to
7 have that term "solely" that we also have in the
8 security research exemption.

9 MS. KARL: And Hacking Policy Council or
10 anyone else, do you know of any examples where this
11 language could cause a chilling effect?

12 MR. GEIGER: Yes. And then I'd like to turn
13 it over to Shayne, who I see had his hand up before
14 me.

15 So, if I recall correctly from the 12th --
16 or, sorry, the 8th rulemaking process, there were two
17 scenarios that we had focused on with "solely." One
18 was academic publishing, so where the language, if it
19 says that it is solely for the purpose of good faith
20 security testing, but an academic also then decides to
21 publish a paper about it, does that go beyond
22 "solely."

23 The second was as part of employment. So,
24 for example, if you're a professional security
25 researcher, and as an example there, there are

1 security researchers that work at companies. I'll
2 give an example, Google Project Zero, that do find
3 vulnerabilities in software that does not belong to
4 Google and they will disclose it to the software owner
5 and they're doing this for the purpose of securing the
6 Internet. And in my opinion, society has benefitted
7 from that. Even when they win a bounty, they don't
8 keep the bounty. They pass it on or donate it to
9 charity, but does that go beyond "solely" since
10 they're doing it for compensation as part of their
11 employment.

12 And if I recall correctly, as part of the
13 8th Triennial process, the Copyright Office clarified
14 that it did not view those things as going beyond
15 "solely," that academic publishing or, you know, as
16 part of your employment would not be -- if those are
17 factors, the Copyright Office does not consider it to
18 be a Section 1201 violation.

19 We felt comfortable with that clarification
20 and that is why we did not pursue any change to
21 "solely" in the Triennial process, and that is why
22 "solely" continues to appear in the language that we
23 had proposed for this 9th Triennial process.

24 MS. KARL: Mr. Longpre?

25 MR. LONGPRE: Maybe I can speak a little bit

1 about the adverse effects part of it. So, in my own
2 lab at MIT, researchers, including myself, were
3 thinking about embarking on trustworthiness research
4 projects to evaluate, in this case, open AI systems.
5 But, after reviewing the terms and even after sending
6 them an email, which was never replied to, we did have
7 lingering concerns about the possibility of legal
8 liability if we were to conduct that research against
9 the terms of service and also if we were to do it and
10 our account was suspended and then we created another
11 account, whether or not that would also maybe
12 engender, you know, more legal liability.

13 And then, in some of our work, co-authors
14 have been red teaming Midjourney looking for ways that
15 the text image model might be unreliable, and in the
16 process of doing that research, they also feared
17 liability, their accounts were suspended, they lost
18 money that they put into it.

19 And from the letter, there were 350-plus
20 people, I'll reiterate, in the research community that
21 have signed on, and one of the bolded parts of that
22 letter, that open letter, is that researchers,
23 independent, the private researchers are experiencing
24 chilling effects when doing good faith academic
25 research.

1 So I just want to say that I think that
2 there's plenty of evidence that this is happening.

3 MS. KARL: All right. Ms. Cohen?

4 MS. COHEN: Yeah. I'll just add that at
5 HackerOne, we partner with good faith researchers.
6 And, you know, we might be retained by a customer, a
7 software developer to look for bias or to do
8 trustworthiness testing and in the course of that
9 testing might identify vulnerabilities with more than
10 just the software that we've been hired to test,
11 namely, an underlying issue with an LLM or a larger
12 player in the market.

13 And there has been some concern about
14 whether or not those should be reported in light of
15 this potential chilling effect and in light of the
16 failure to protect individuals who are doing that
17 underlying research. Again, not with the customer who
18 has retained us to actually do that testing but the
19 underlying LLM.

20 MS. KARL: All right. We're going to go to
21 Mr. Reed and Mr. Englund very briefly before moving to
22 the next question.

23 MR. REED: Thank you. I'm struggling in
24 large part because most of my membership and others,
25 we kind of align with the proponents, but I just heard

1 Shayne talk about that no email was answered. But
2 Harley Geiger has Microsoft as a member of the
3 Advisory Committee for Hackers. Charlie Snyder from
4 Google is on the Alliance.

5 Heck, you could've emailed me. I mean, we
6 work really closely with Microsoft on Health. And so
7 I'm a little worried that we're being asked to add an
8 exception where a phone or an email would be possible,
9 because Mr. Geiger is a leading expert on this. I
10 remember him when we all worked on the Hill. It's an
11 email to him and he's probably going to be able to
12 reach to Microsoft, and I'm disappointed that they
13 didn't email you back. You've got great advisors at
14 MIT.

15 So I'm trying to figure out how to be with
16 you, but I'm hearing solutions that are asking for the
17 Copyright Office to move on 1201 with problems that we
18 could probably solve within our industry through
19 better communication. So my apologies if you didn't
20 get the support you needed to do that research and I,
21 as a member of the industry, can try to do better.
22 But we're out here and Mr. Geiger's got some of those
23 experts on his Advisory Committee. So let's figure
24 out if we can solve this faster than government.

25 MR. ENGLUND: And I'd like to respond just

1 briefly to Mr. Longpre's comments a few minutes ago.
2 I think I heard him say that a project he was working
3 on was discontinued because of a concern about terms
4 of service violations and he went on to describe some
5 other concerns as well.

6 But, if a concern over violating a service's
7 terms is killing projects, nothing else is going to
8 matter because the Office can't immunize researchers
9 from terms of service violations and contract
10 liability.

11 And similarly, Mr. Longpre referred to the
12 open letter that's attached to the academic
13 researchers' comments. I searched that letter to try
14 to find any reference to circumvention or Section
15 1201. I just couldn't find it. It was all about
16 terms of service violations. And at the risk of
17 repeating myself, the Office just can't immunize
18 researchers from contract liability for violating
19 terms of service.

20 MS. KARL: Thank you.

21 We're going to have to move on.

22 MR. LONGPRE: Can I respond at some point,
23 Ms. Karl, to those points?

24 MS. KARL: We really have to move on. And
25 if we have time or if you can work it in, that would

1 be great.

2 Yeah. So we wanted to actually go back to
3 something that was raised earlier. The Office is
4 currently conducting a study regarding the copyright
5 issues raised by generative artificial intelligence.
6 Because it is rapidly evolving and because the
7 Office's study may touch on related issues, should the
8 Office wait three years before opining on an AI-
9 related research exemption?

10 Ms. Elazari?

11 DR. ELAZARI: Yeah. I'll be brief. I think
12 it's important to note that there are perhaps
13 questions that are open on the copyrightability of AI,
14 but there is overwhelming support, including in the
15 Executive Order, including by CISA, by NIST, and by
16 other agencies, including those on the line here, that
17 testing of AI is important. That red teaming of AI and
18 specifically third-party red teaming, which is, by the
19 way, required by law in different states, is
20 beneficially -- you know, it's a beneficial social
21 activity.

22 So I think we need to distinguish the fact
23 that while there might be an open question on
24 copyrightability of AI, the question on whether
25 testing of AI and the importance of finding those

1 unintended consequences that the White House and
2 others are concerned about, those are well
3 established, right? So I think this is just something
4 that I wanted to raise.

5 And, in fact, we have seen already not just
6 the establishment of an AI Safety Institute by the
7 Department of Commerce but the creation of a specific
8 working group on the importance of AI retaining. So
9 we know already that this type of testing and work
10 that is being done by third-party researchers is not
11 just acknowledged. It's about to be required in
12 certain segments of the market, and, therefore, I
13 think that is an important distinction that I would
14 like to draw, suggesting that we should not wait
15 because we're not waiting on asking those important
16 communities to inform us with their testings and with
17 their findings.

18 MS. KARL: Thank you.

19 Mr. Geiger?

20 MR. GEIGER: Yes. So you asked whether we
21 should wait three years for this to come around again.
22 I would suggest no, that Section 1201(a)(1)(c) asks
23 for a preponderance of evidence whether this will
24 likely adversely affect non-infringing uses in the
25 three-year period following this proceeding.

1 And as you can see from the record,
2 particularly that built up by Mr. Longpre, you are
3 hearing about specific adverse effects. You're
4 hearing about a community of hundreds and hundreds of
5 researchers that are worried about adverse effects.

6 I would suggest that the preponderance of
7 evidence demonstrates that it is likely to have
8 additional adverse effects in the subsequent
9 three-year period following this proceeding.

10 In addition, the computer programs - or I
11 should say, the study that the Copyright Office is
12 engaging in regarding the copyrightability of AI, I'll
13 just note that, again, what we are focused on are the
14 code for the user interface, code for the API, and
15 software code that drives the algorithm. I would be,
16 frankly, shocked if the outcome of the Copyright
17 Office's study is that the code for those computer
18 programs are not protected works.

19 And then, lastly, I would just reiterate
20 that we are not looking for immunization for all
21 liability from deviating from terms of service. That
22 would be inappropriate and that is clearly not what
23 Section 1201 can or should do. We are only asking the
24 Copyright Office to provide protection from liability
25 under Section 1201. And it would not be a sufficient

1 alternative to ask researchers to work with every AI
2 system operator for every active research.

3 MS. KARL: All right. Mr. Longpre?

4 MR. LONGPRE: Yeah. To answer your
5 question, Ms. Karl, about the timeliness or can we
6 wait three years, I think the answer is no because of
7 how essential and critical this research is right now,
8 as echoed by the community, but also because the
9 alternatives aren't really viable.

10 So I think Mr. Reed mentioned there needs to
11 be better communication between researchers and
12 companies, but as Mr. Geiger just said, it's virtually
13 untenable for thousands of researchers investigating
14 general purpose models with so many different uses,
15 from law, medicine, education. Children are using
16 this in schools and outside of schools. And there are
17 so many different places that need to be investigated
18 that we know these companies are understaffed and have
19 maybe one or two people looking at these applications
20 or answering these emails, and it's virtually
21 impossible for this communication to be alternative,
22 in my view, to this exemption.

23 And just to address the comments about the
24 letter that Mr. Englund brought up, it specifically
25 mentions in that very short letter fear of legal

1 reprisal and chilling effects, and that's the thing
2 we're talking about. We don't mention specific
3 statutes in there because it's a community letter, but
4 those are, as we understand it, the primary concern of
5 legal liability.

6 MS. KARL: All right. Mr. Englund, and then
7 we're going to move to the next question.

8 MR. ENGLUND: Yeah. So I agree with the
9 premise of your question that it is premature to adopt
10 an exemption at this time and refer you to our written
11 comments, which address that at greater length.

12 But I believe that to be true for several
13 reasons. First, the record here is very incomplete,
14 and I don't think that the Office is in a position on
15 this record to make a judgment that an exemption is
16 appropriate because of likely adverse effects over the
17 next three years.

18 But more generally, the Office's AI study is
19 very wide-ranging. It is not focused simply on
20 copyrightability issues and, as described in our
21 written comments, does potentially implicate issues
22 that are relevant here. But NIST and others are
23 studying the red teaming issues. This is a very
24 dynamic environment. The issues are all novel and it
25 doesn't seem like the time to be acting on an

1 incomplete record in the absence of knowledge on
2 exactly how things are going to shape up.

3 MS. KARL: All right. Thank you.

4 For opponents, several reply comments, along
5 with the exemption language within those comments,
6 suggests that the exemption could be used to
7 investigate the extent to which AI models reproduce
8 copyrighted material.

9 If the Office were to grant the proposed
10 exemption that encompassed research into infringement,
11 would that be something that you would find desirable?

12 MR. GEIGER: So I'll start that. That is
13 the language that we had proposed. Yes, that is
14 indeed within the language that we are proposing.
15 What we are describing as AI trustworthiness has a
16 specific definition. It's actually relatively clear.
17 These definitions for trustworthiness include several
18 concepts. They include bias. They include
19 resiliency. They also include validity and
20 reliability, and infringement would be a type of
21 reliability harm.

22 So, under the definitions that NIST uses, as
23 well as other international standards, specifically,
24 ISO IECT 5723, I know that's a mouthful, but that is
25 the origin of the trustworthiness definition.

1 Reliability is a goal for overall correctness of an AI
2 system operating under the conditions of expected use.

3 So, presuming that the AI system is not
4 designed to produce infringing material, then an AI
5 system that does produce infringing material is not
6 operating correctly under conditions of expected use.
7 And, therefore, a researcher that is able to show that
8 an AI system can produce infringing material against
9 its intent has identified the trustworthiness problem
10 of reliability, so, yes, and our language encompasses
11 that type of research.

12 MS. KERN: Thank you. Yeah. And just
13 really quickly, Mr. Geiger, is the reason that you
14 changed your language, proposed exemption language,
15 between the initial comment and reply comment from
16 alignment to trustworthiness for the reason that you
17 just stated?

18 MR. GEIGER: No, not that particular reason.
19 The reason is because I used alignment -- and to be
20 clear, to my knowledge, that's the only thing that has
21 changed from our original language, is switching out
22 the word "alignment" for "trustworthiness." It is
23 simply because, although "alignment" does have a
24 definition that generally means keeping AI consistent
25 with societal norms, and so it could work, but there

1 was just a stronger body of evidence and general
2 acceptance by the community, the standards community,
3 around the word "trustworthiness" instead.

4 So alignment and trustworthiness could both
5 work, but because of the use of trustworthiness in
6 NIST's AI risk management framework, NIST -- other AI
7 trustworthiness work, as well as the ISO language
8 around trustworthiness we thought that that was the
9 clearer term to use here.

10 MS. KERN: Thank you. Just wanted to
11 clarify.

12 And, Mr. Ayers, I'll give you 15 seconds
13 because I want to move on to my colleague, Kevin. I
14 know he has a question he wants to ask.

15 MR. AYERS: Sure. No problem.

16 I'm just noting that I think it's a tough
17 call because the concern would be a cost/benefit
18 analysis. Is there enough benefit in the possibility
19 of infringing material being identified down the road
20 that it more than makes up for the risk to all the
21 rest of the material that, for instance, might be
22 exposed because of a circumvented Blu-Ray player?

23 MS. KERN: Thank you.

24 And, Kevin, the floor is yours.

25 MR. LI: Thank you.

1 I have a question about how the current
2 request interacts with the existing exemption on
3 security research. And in particular, for proponents,
4 it would be very helpful for you to discuss any ways
5 in which the current exemption is insufficient for the
6 purposes for which you're hoping to circumvent TPMs.
7 And it would also be helpful for opponents to discuss
8 if there are any ways in which that seems overbroad,
9 but let's start with proponents.

10 Mr. Geiger?

11 MR. GEIGER: Yes. So I would actually like
12 to cite the Department of Justice letter here. I
13 defer very much to the enforcers of our intellectual
14 property laws where they say, "While the existing
15 exemption for computer security research covers many
16 types of research focused on the security and
17 integrity of AI models, we recognize that it may not
18 be sufficiently broad in its current form to exempt
19 research that falls outside of security concerns."

20 The Department of Justice agrees that an
21 exemption focused on security is possibly not going to
22 cover non-security harms present in AI, such as bias,
23 discrimination, and other trustworthiness issues that
24 we've described here.

25 MR. LI: Are there any other proponents that

1 would like to speak to this issue?

2 DR. ELAZARI: Yeah. This is Dr. Amit
3 Elazari from OpenPolicy. I would just like to, again,
4 echo Harley's note and then the underlying statement
5 that the Department of Justice has provided.

6 I think, you know, while there is ambiguity,
7 it's very clear that there are a whole set of
8 unintended consequences that can be stemming from AI
9 systems. Many of them are cited again in the
10 Executive Order from bias to discrimination to
11 reliability and trustworthiness. And as the
12 Department of Justice suggested, because of this broad
13 set of unintended consequences, it is important that
14 we create this exemption.

15 MR. LI: Are there any opponents that would
16 like to speak to this question?

17 (simultaneous discussion)

18 MR. TAYLOR: Yeah. I would just say that
19 the problem with what I see is, one, we don't have a
20 record, but the distinction is, in the current
21 security research exemption, is what constitutes harm.
22 And, obviously, when we had the word "harm" on the
23 records that were created previously, we had examples
24 of what that harm was, and, here, we don't have a
25 record that distinguishes the harm.

1 And, you know, the Copyright Office will
2 certainly give sufficient weight to the Department of
3 Justice, but there's nothing that says the Department
4 of Justice dictates what the exemptions will be. And
5 so the Copyright Office will follow the law and will
6 recommend an exemption based on the record, and there
7 is nothing here.

8 MR. LI: And, Mr. Englund, if you have
9 something brief you'd like to say.

10 MR. ENGLUND: Yeah. So the original
11 petition that was filed and led to this class referred
12 specifically to security in generative AI, and,
13 obviously, we've gone well beyond that now.

14 You know, I think that the current exemption
15 speaks for itself. It's about security. And so, to
16 the extent we've gone beyond security, yeah,
17 presumably, the current 16 exemption doesn't cover it.
18 I think that's a problem and an illustration of the
19 breadth of the class that's being proposed.

20 But trivially enough,
21 I do want to agree with Mr. Taylor's comment here that
22 I suspect that, you know, CSIP's jurisdiction includes
23 things like security. And as I read the letter, it
24 was very much focused on security. It is not clear to
25 me that they appreciated the full breadth of the scope

1 or that we might be talking about providing access to
2 the kinds of creative content that, based on today's
3 discussions, sounds like might be potentially in the
4 cards given this exemption.

5 In any event, I certainly agree with the
6 proposition the Office needs to apply copyright law on
7 its own and is not to defer to the opinions of any
8 other administrative agency.

9 MR. LI: Thank you, Mr. Englund.

10 And very briefly, Mr. Geiger?

11 MR. GEIGER: Yes. I'd like to just respond
12 to what Mr. Englund said about the Department of
13 Justice's letter being security-focused and just to
14 note that page 4 of that letter states, "CSIP believes
15 that good faith research on potentially harmful
16 outputs of AI and similar algorithmic systems should
17 be exempted from the DMCA circumvention provisions."

18 Just flat out, that's not talking about
19 security. It is, in fact, recognizing that security
20 is alone insufficient to cover non-security harms
21 under Section 1201 and they have flatly recommended
22 that they be exempted.

23 MR. LI: Thank you, Mr. Geiger.

24 I'm going to pass it back to the Copyright
25 Office. Ms. Karl.

1 MS. KARL: Thank you, Kevin. Thank you.

2 This has been referenced a couple of times.
3 So we're interested to know, is industry
4 self-regulation sufficient to manage data provenance
5 and bias issues in AI research, both generative and
6 overall? If so, then why? And if not, is external
7 research into these questions needed?

8 Mr. Longpre?

9 MR. LONGPRE: Sorry. Can you repeat the
10 question? Is the industry able to self-regulate data
11 provenance and bias, is that what you're asking?

12 MS. KARL: Yeah. To self-regulate the
13 research into these questions.

14 MR. LONGPRE: I see.

15 MS. KARL: Do you want to answer, or do you
16 want to let Mr. Reed?

17 MR. LONGPRE: I'd be happy to answer
18 quickly.

19 MS. KARL: Okay. Great.

20 MR. LONGPRE: I think that, very broadly
21 speaking, the community does not think that it is able
22 to self-regulate in a way that's in the public
23 interest or able to give consent to the broad set of
24 good faith researchers that want to do analysis into
25 many different aspects of these systems without them

1 incurring chilling effects because, in some cases, the
2 companies don't want there to be investigations into
3 their vulnerabilities or the ways their models can
4 produce bias. In other cases, they're understaffed
5 and so they just don't respond to researchers, but in
6 either case, I think the answer is no.

7 MS. KARL: Mr. Reed? And just to clarify,
8 you know, we're kind of interested in voluntary
9 agreement in terms of self-regulation.

10 MR. REED: Right. So I'm actually going to
11 side, so to speak, with the proponents in the sense
12 that I don't think pure industry self-regulation is
13 sufficient. The good news there is, is that's not
14 what's happening.

15 I have the good fortune to work with Health
16 and Human Services and with their Office of Civil
17 Rights in which they're engaging directly with
18 agencies. The Food and Drug Administration similarly
19 had its kitchen cabinet around AI. And so each agency
20 has now been working with industry but in a
21 quasi-self-regulatory way in the sense that it is
22 government asking the questions and engaging
23 researchers.

24 So I think it's a misnomer to suggest that
25 industry could do this on its own and that we live on

1 some magical island. However, I think that right now
2 what we're seeing is every agency is working directly
3 with providers and especially with foundational model
4 developers to work on these questions. So, no, it's
5 not pure industry self-regulation and that's not
6 what's happening today.

7 MS. KARL: Thank you.

8 Mr. Geiger?

9 MR. GEIGER: Yes. I would argue that this
10 exemption will actually very much aid the industry in
11 self-regulating. So industry itself, as noted, is not
12 able to realistically identify these algorithmic flaws
13 by themselves. AI systems are presently
14 decentralizing rapidly. We're seeing them become more
15 and more accessible by more and more industry players,
16 some of whom we're not sure where they are geolocated.
17 And it is untenable to expect that research would have
18 to negotiate with each one of these actors, let alone
19 what that would do to the independence of the research
20 for each negotiation.

21 In addition, even for large organizations,
22 well-resourced organizations that do try to take steps
23 to ensure the trustworthiness of their AI models, you
24 hear these systems described as a black box. The fact
25 is that many even large organizations that run AI

1 systems don't often know what their AI system is going
2 to do. And so it is not realistic to expect them to
3 identify algometric flaws that they don't even know
4 are there.

5 So enabling researchers to, on an
6 independent basis, without fear of liability under
7 Section 1201, find these flaws and submit them to
8 industry and talk about them in academic conferences
9 will help the industry self-regulate. Really, in the
10 end, what we're asking for here is deregulation.

11 MS. KARL: Mr. Harguess?

12 DR. HARGUESS: Yeah. I'll agree with pretty
13 much everything that's been said so far on this. I'll
14 double down on, you know, large organizations that
15 have the resources, you know, they can stand up their
16 own red teams. You know, sometimes they're putting
17 language in that says, you know, you can't red team
18 our model. We're doing that on our side of the fence.

19 You know, I think we want to enable a rich,
20 you know, landscape of researchers that can do this,
21 you know, independently. And further, we know that
22 some governance is going to come down that requires
23 independent red teaming and independent testing.

24 MS. KARL: Ms. Elazari? And then we will
25 move to the next question.

1 DR. ELAZARI: Yeah. I agree. I think, you
2 know, as a matter of policy, it's well established
3 that this type of testing activity is something that
4 is desirable. In fact, we already have laws, and I
5 refer you to a law that is in the State of New York
6 that requires third-party audits on certain systems in
7 the context of trustworthiness.

8 And on the other issue, I'm also the
9 co-founder of Disclose IO, which is a prominent set of
10 private ordering contracts that are being used in
11 order to allow safe harbor activity using contracts.

12 And I very much agree with Mr. Geiger and
13 others that, you know, thinking that private industry
14 would go and roll out those private consents is very
15 much unrealistic even in security. Even though we
16 have all this progress, including frameworks like
17 Disclose IO being promoted by CISA and others and
18 required by federal agencies, we still don't have
19 broad adoption of that contractual language.
20 Certainly, to the case we have it, it's focused on
21 security. So we certainly need this action from the
22 Copyright Office in order to promote this desirable
23 activity, as Mr. Geiger mentioned.

24 MR. LONGPRE: I can add really quickly that
25 the current reality is that very few maybe elite labs

1 and institutions, organizations usually that have
2 connections with these well-resourced companies are
3 the ones that are given special permissions to do this
4 research, and the much broader community is left
5 usually with some form of chilling effects or
6 uncertainty or not hearing back and, as a result,
7 isn't doing that research when, you know, there are a
8 hundred million plus people using these services from
9 across the world and two years ago that was zero.
10 It's the fastest growing, and so there are so many
11 vulnerabilities that require all the different
12 communities to participate, and self-regulation isn't
13 getting us there.

14 MS. KARL: Okay. Thanks.

15 Melinda?

16 MS. KERN: Thank you very much.

17 So I just had a quick question. We're
18 moving, Brandy got to it a little bit, to the
19 non-infringing use section and then we're going to go
20 to adverse effects. And, unfortunately, because of
21 time, we might have to make these responses brief.

22 But the Office is not aware of at least at
23 the moment any case or legal authority ruling that the
24 actions under this specific proposed exemption are
25 non-infringing. In fact, we said some statements

1 about fair use in our previous triennial rulemaking,
2 but are proponents or opponents aware of any authority
3 on the question of non-infringing uses, such as fair
4 use or Section 117, for good faith AI security
5 research? And that, again, is for both proponents and
6 opponents.

7 Mr. Englund?

8 MR. ENGLUND: So I'm not aware of any
9 authority, but it is conspicuous in the record here
10 how little attention there has been to the question of
11 whether uses are infringing or not infringing.

12 Since you raised 117, I'll note that it's
13 not clear to me how that would apply to most of the
14 use cases that have been talked about today. No user
15 owns a copy of ChatGPT, so that just seems totally
16 irrelevant, meaning that it boils down to fair use.
17 And, here, I think the fair use analysis hasn't really
18 been talked about much in the comments, is very
19 different from the kind of analysis that has often
20 been possible when the Office granted exemptions.

21 So, here, the proposed exemption is not
22 limited to noncommercial users. In fact, I understand
23 that a number of the witnesses today are
24 representatives of commercial entities, so the Warhol
25 case told us we need to take into account the

1 commerciality of the use.

2 And in terms of the nature of the use, we've
3 heard today that this exemption would potentially give
4 the ability to access creative content and be able to
5 use it however it might be available once the access
6 to a system has been circumvented, and so I think the
7 first factor is problematic here.

8 Similarly, factor two, if we're talking
9 about circumventing TPMs on DVD players and streaming
10 services and video games, we're potentially talking
11 about creative works. I don't know if we have enough
12 of a record to judge how much copyrighted works need
13 to be copied for these purposes. The proponents just
14 haven't told us very much about that.

15 And similarly, for the fourth factor, it
16 seems like market harm is a possibility if we're
17 talking about exposing creative works at least.

18 MS. KERN: Thank you.

19 Mr. Geiger?

20 MR. GEIGER: So I would strongly suggest
21 that the fair use analysis is identical to the
22 analysis that the Office conducted in past triennial
23 rulemakings for security testing. We have no interest
24 in and do not want to see this exemption used for
25 infringement. This exemption is directed at fair use.

1 AI trustworthiness research is fair use. It
2 contributes to the advancement of computer science,
3 and it leads to the production of new creative works.

4 I'll note that within the 8th Triennial
5 proceeding the Register of Copyrights said that in
6 prior rulemakings the Office has consistently found
7 that exemptions to allow non-infringing analysis of
8 computer programs are likely to promote the
9 availability of copyrighted works.

10 I'll also point out that the language that
11 we are proposing specifically states that the results
12 of the research would not be used or maintained in a
13 manner that infringes on copyright. We have, in fact,
14 tried to craft our exemption request in such a way
15 that infringement would fall out of the exemption.
16 And we believe this to be a fair use activity, again,
17 very much in line with the analysis that occurred
18 under security testing.

19 If we need to go into a deeper analysis of
20 fair use, I think that that would be possible. But,
21 again, I feel relatively confident in this conclusion
22 that such research is fair use.

23 MS. KERN: Thank you. And I should've
24 clarified. The question was more directed towards
25 authority, but I will let you go ahead, Mr. Taylor and

1 then Mr. Reed.

2 MR. TAYLOR: Yeah. Yes. I would just
3 simply say that this rulemaking with every exemption
4 requires an evidentiary record and there can be
5 nothing that's assumed based on any other exemption
6 that appears to be similar. And we've evolved
7 significantly since 2001 and that evolution has only
8 gone through renewals. But we cannot create an
9 exemption based on, oh, yeah, this smells, sounds like
10 something that we've already created. So, with that,
11 I don't think there's a record here.

12 MS. KERN: Thank you.

13 Mr. Reed?

14 MR. REED: Yes. I'm doing the dangerous
15 thing of asking a question I don't know the answer to
16 here. But borrowing from the previous session that we
17 just had, Mr. Geiger's point about you can't use the
18 work done from research in a way that would be
19 profitable. But what if I were to red team you and
20 pull that data down on the way that you implement your
21 LLM and I were to train my LLM with it if I was a
22 competitor?

23 Let's say I'm a pentesting company and you
24 do great work and I look at how you're implementing
25 it. I want to research your system for bias. I can

1 use that to train because, as we just discussed, it's
2 unclear where that use is. And for those of you who
3 were in the previous session, there was a major
4 conversation about can I use someone else's data in
5 such a manner.

6 So I'd argue, like the previous commenter,
7 about the body of evidence right now. And I'd need to
8 think really hard about does this open the door for
9 someone not exactly stealing your copyrighted material
10 but rather training their own LLM under the auspices
11 of research and checking because it's not, in fact,
12 taking your copyrighted material. Thank you.

13 MS. KERN: Thank you.

14 And for timing purposes, Mr. Geiger, since
15 you already responded, I'll give you 15 seconds.

16 MR. GEIGER: Thank you. I would note that
17 our comments did, in fact, claim that this was fair
18 use. This is on the record. It is at the bottom of
19 page 6 of our reply comments. There's also absolutely
20 no evidence that the proposed exemption would result
21 in increased copyright infringement or piracy.

22 With regard to the example that was just
23 supplied by Mr. Reed, I would argue that that is
24 likely out of the scope of solely for purposes of good
25 faith AI trustworthiness research. If you're building

1 a commercial product with the results, that seems that
2 we've gotten farther afield from academic publishing
3 or if this is part of your employment and you are a
4 professional researcher.

5 So, no, I don't think that qualifies. And
6 you can make the same arguments with existing
7 exemptions, such as security testing. Thank you.

8 MR. REED: Thanks.

9 MS. KARL: Yes. I have a question for
10 proponents. Are there white papers or other manuals
11 collecting techniques for the kind of security
12 research that you hope to engage in that describe the
13 variety of techniques that are used for this kind of
14 research? Mr. Longpre? I'm sorry.

15 MR. LONGPRE: Ilona, do you want to go
16 ahead?

17 MS. COHEN: No. I'm sorry. I didn't see
18 that you raised your hand. Sure. Yeah. HackerOne
19 has published a number of different guides for AI
20 trustworthiness and red teaming, which are made
21 available on our website and are drawn from our
22 multitude of experiences doing this red teaming for
23 customers.

24 MS. KARL: Thank you.

25 Mr. Longpre?

1 MR. LONGPRE: We'd be happy to provide lots
2 of resources related to that. There are many papers.

3 MS. KARL: Mr. Geiger?

4 MR. GEIGER: Yeah. Just to say that there
5 are, indeed, numerous white papers describing AI red
6 teaming and testing techniques. And if you'd like to
7 know more, we can certainly provide some of them to
8 you as well.

9 MS. KARL: Mr. Harguess?

10 DR. HARGUESS: A very similar response.
11 We've produced webinars, you know, several technical
12 documents back at my time with MITRE, same thing,
13 MITRE ATLAS has been stood up. That talks a lot about
14 some of the things we're discussing here, so happy to
15 provide some materials.

16 MS. KARL: Ms. Elazari, you raised your
17 hand?

18 DR. ELAZARI: Yeah. I would just echo there
19 are thousands of papers, including those produced by
20 the hundreds of scholars that -- you know, the letter
21 that Shayne referred to that describe the discipline
22 of AI auditing and algorithmic auditing, and that's a
23 record or that's information we're happy to provide.

24 MS. KARL: And Mr. Reed? Oh, you are muted.

25 MR. REED: Somebody had to do it. Sorry I

1 drew the short straw. Our Connected Health Initiative
2 has a very comprehensive trustworthiness guideline
3 that also provides insight into each level of
4 responsibility. This is particular for the healthcare
5 industry, but it's appropriate very broadly, and we'll
6 make sure to submit that for the record.

7 MS. KARL: I'm passing it back to Melinda.

8 MS. KERN: Thank you.

9 So I wanted to get into alternatives a
10 little bit. So this is a question for the supporters,
11 proponents, and then I'll do a follow-up for the
12 opponents of the proposed exemption.

13 So HackerOne asserts that while good
14 research, access, and bias bounty programs are
15 available for identifying things like bias, these
16 programs are often limited in availability and scope.
17 And I also believe the Department of Defense had a bug
18 bounty program for AI bias that I believe ended in
19 February 2024.

20 Could you please speak a little bit more to
21 why these are allegedly unreasonable alternatives and
22 what the proposed exemption covers? And you can do
23 that last part briefly because I know we touched on it
24 a little bit, but please go ahead.

25 MS. COHEN: Sorry. I didn't hear. Why what

1 is unreasonable?

2 MS. KERN: Why things like research access
3 and bug bounty programs that are currently available
4 are limited in availability and scope and why they
5 aren't reasonable.

6 MS. COHEN: Well, as you mentioned, we did
7 discuss this. It's generally that the current scope
8 and the protections offered are for security testing
9 primarily, and so the expansion of testing for
10 trustworthiness, for bias, for discrimination is
11 necessary in order to be able to cover the scope of
12 programs that we do, including for the Department of
13 Defense and other government customers.

14 MS. KERN: Thank you.

15 Mr. Longpre?

16 MR. LONGPRE: If you're talking about the
17 company programs, they are very limited in scope.
18 They self-select who gets to opt in and do that
19 research by an application pool. We know many, many
20 top tier researchers that never got their applications
21 accepted or heard back from any of them. Not all
22 companies even have these programs and so it turns
23 into like a very small set of researchers that get to
24 do this that don't have necessarily the independence
25 that third-party independent good faith researchers

1 would have.

2 MR. KERN: Mr. Geiger?

3 MR. GEIGER: So I would just say that
4 research access programs and bias bounties are very
5 good practices that some elements of the industry are
6 engaging in right now. However, it is not something
7 that is industry-wide. And as noted, those types of
8 programs generally place rules around how the TMPs can
9 be circumvented, so which circumstances, which
10 methodologies, which assets, and then they have terms
11 regarding disclosure.

12 In addition, they're not made available to
13 the entire community of good faith researchers. So
14 there are helpful programs, but they are not an
15 equivalent.

16 I'll also note that what we are seeing with
17 AI is that AI instances are being licensed in other
18 places. So an AI model owner will create the model
19 and may provide, you know, an interface directly to do
20 there, but also other applications can take an
21 instance of that and license and instance of it.

22 So let's just say an eCommerce platform, for
23 example, licenses an instance of a generative AI model
24 on the eCommerce platform. In that circumstance, you
25 would have to be looking -- you know, the eCommerce

1 program itself may not have its own bias bounty or its
2 own research access program, but the fact that it is
3 using the generative AI model in its instance can
4 create important new avenues of research.

5 Now the fact that it's in an instance could
6 yield unique results. And, of course, there is the
7 API, the technical environment within which that
8 instance sits. So these researcher access programs
9 are very helpful but absolutely not an equivalent for
10 enabling independent, good faith trustworthiness
11 research. Thanks.

12 MS. KERN: Thank you.

13 Ms. Elazari?

14 DR. ELAZARI: Yeah. I just wanted to double
15 down on the issue of the inclusivity of the program.
16 So, as Mr. Geiger and Mr. Shayne mentioned, those
17 programs are very, very, very limited. They are not
18 currently widely adopted, and they are really open to
19 a selected few. And in the context of AI
20 trustworthiness research, there is an important
21 element of the diversity and the skillset that is
22 rooted in diversity of the auditors themselves.

23 So that is to suggest this type of scoping
24 is, you know, just doubling down on the comments being
25 made. It's especially problematic because it's the

1 type of diversity and diverse backgrounds of testers
2 that we want to enable in the context of bias and
3 discrimination research.

4 MS. KERN: Thank you.

5 Mr. Reed?

6 MR. REED: So I want to address briefly the
7 bias bounty and the concept that if some is good, more
8 is better. That may not be true. The largest problem
9 that we have with the concept of bias bounties and why
10 a little more restriction is the norm is that bias is
11 not bullion. It isn't like pentesting. Did I get
12 root or not? Was I able to get access to your entire
13 machine?

14 Bias is obviously something in which both
15 cultures, areas, terms are different. And, therefore,
16 one of the things that everyone has been concerned
17 about in the bias bounty area is do we gameify it so
18 much that it becomes not as clear.

19 In pentesting and red teaming on the
20 security front, what I'm normally trying to do is can
21 I have access to something I am not supposed to have
22 access to. In bias testing, I'm looking for something
23 that is, is this result what I would expect as a
24 person of this culture or of this norm, or in the
25 healthcare implementation, it gets very interesting

1 because of how physicians provide records that go into
2 the electronic health record.

3 So I think that we shouldn't just assume
4 magically that if there's some difficulty in doing bug
5 bounties - I'm sorry, bias bounties it's all because,
6 gosh, nobody thought of this and we just can't make it
7 work. There are practical reasons.

8 And then, finally, I would kind of ground
9 this back into why we're here at a 1201 procedure,
10 which is the Copyright Office needs to see a body of
11 work that is significant before making significant
12 changes. We have a triennial process that starts with
13 a presumption that the copyright of the person that
14 developed the material has some right of exclusivity,
15 some right to prevent, and the opportunities that have
16 been created through security and other exemptions are
17 very specific for a reason, because it starts with the
18 position that the copyright holder has a certain
19 amount of right to protect their material.

20 So let's be cautious by saying more is
21 always better on bias bounties. And then, second,
22 let's reground this in the conversation, which is the
23 purpose of the Copyright Office is to ensure that
24 those who have that right are able to exercise it
25 appropriately. And until we see a larger body of

1 work, I'm not sure we can jump to the conclusion that
2 we should just open all the doors.

3 MS. KERN: Thank you.

4 Did any other opponents have any comments on
5 whether they thought those programs or research access
6 were reasonable alternatives?

7 (No response.)

8 MS. KERN: Okay. I'm not seeing any hands
9 up, so I will move on to our next question, and this
10 kind of goes back a little bit, and I'm not sure if
11 everyone touched on it before. But, besides those
12 discussed in the comments and what's already been
13 mentioned during this hearing here, are these the only
14 TPMs that need to be mentioned or are there some that
15 haven't even been mentioned by everyone? And that's a
16 question for the proponents and then also for the
17 opponents. Are you aware of any other TPMs that would
18 need to be accessed in the use cases that proponents
19 have described that you haven't already mentioned?
20 Just because we're running low on time, I just want to
21 caveat the responses.

22 Mr. Longpre?

23 MR. LONGPRE: In our comment, I believe we
24 have a list of 10 that breaks some of them down more
25 finely that might be considered. I can look that up

1 now.

2 MS. KERN: And Mr. Englund?

3 MR. ENGLUND: It's a very difficult question
4 to answer because the record here is very thin, but if
5 you accept that this exemption would allow
6 circumvention on TPMs of video games that have AI
7 features, that also implicates the TPMs in video games
8 consoles because the TPMs on consoles and games
9 interoperate to provide a secure operating system for
10 playing games.

11 And there have been plenty of record in
12 other proceedings about the security and piracy issues
13 associated with hacking of consoles, but that could
14 somehow be implicated here depending on how broad you
15 think the exemption is.

16 MS. KERN: Thank you.

17 And Mr. Geiger?

18 MR. GEIGER: Yes. I would just point you to
19 the comments from the Joint Academic Researchers,
20 which lists numerous technological protection measures
21 that may be implicated. There are probably others
22 that, you know, we could try to find, but there are a
23 variety.

24 I would just note also on the continued
25 referring to the record as being thin, you know, we

1 have provided very specific language, provided very
2 specific TPMs. We've provided specific instances of
3 adverse effects, as well as noting a broader community
4 of fear of adverse effects of Section 1201.

5 I would argue that the record is actually
6 very robust in terms of a standard of preponderance of
7 the evidence that there will likely be additional
8 adverse effects in the coming three years.

9 MS. KERN: Thank you.

10 And then I just had one question, and then
11 I'm going to pass it to Mr. Li. We're talking about
12 TPMs, and I know that both Hacking Policy Council and
13 Joint Academic Researchers really describe the terms
14 of account or terms of service, which is outside the
15 realm of copyright law and outside the realm of this
16 rulemaking, as others have mentioned before, as
17 prohibiting the activities that the proposed exemption
18 seeks to permit.

19 Would you mind just very briefly discussing
20 why you believe the terms of account or terms of
21 service are not actually the cause of the adverse
22 effects for the security researchers?

23 Mr. Geiger, go ahead.

24 MR. GEIGER: I'll turn it over to Shayne
25 very quickly. The terms of service in many ways are

1 beside the point. The terms of service are not the
2 technological protection measure. When we describe
3 TPMS, we are describing account requirements. We are
4 describing rate limits. We're describing guardrails.

5 The terms of service really come into play
6 when you lose your accounts. They're the reason
7 sometimes that an individual loses their account.
8 However, you can have your account suspended for any
9 reason even if you're not violating terms of service.
10 So this is really not about terms of service and
11 changing terms of service. This is about removing
12 liability under Section 1201 for good faith AI
13 trustworthiness when you are circumventing TPMS that
14 include account suspension.

15 And I would just reiterate also that we
16 fully understand and our language does not contemplate
17 prohibiting AI system operators from having terms of
18 service to prevent behavior that they find
19 undesirable, including good faith research. A Section
20 1201 exemption does not prevent that. A Section 1201
21 exemption merely helps reduce the adverse effect that
22 is created by liability under Section 1201. So, to
23 the extent that operators still want to be able to
24 prevent undesirable conduct taking place on their
25 platforms, they still retain every right to do so.

1 MS. KERN: Thank you.

2 And just before I let you go, Mr. Taylor,
3 I'm just going to set a roadmap because we have very
4 little time left. I'm going to have you respond, Mr.
5 Taylor, pass it to Mr. Li, and then we'll do closing
6 remarks. Thank you.

7 MR. TAYLOR: Yeah. I just wanted to point
8 out to the Copyright Office that the Register back in
9 the 2003 opinion did look at end user license
10 agreements and said that end user license agreements
11 were not subject to 1201 and that they were a separate
12 contractual violation. So, if you go back, I think it
13 starts with a discussion is around page 149 of the
14 2003 opinion. And sorry I didn't cite it earlier in
15 our opposition.

16 MR. LI: Thank you very much, Mr. Taylor.

17 And thank you for the pass, Ms. Kern.

18 I'd like to zoom out very slightly here and
19 discuss adverse effects more broadly, and this is
20 principally for proponents.

21 If you can discuss, you know, the need for
22 trustworthiness research now and over the next three
23 years prior to the next Triennial and, you know, if
24 there are any concerns that you have about if there
25 isn't an exemption, you know, what kinds of harms

1 might result?

2 MR. LONGPRE: I'm happy to speak to that a
3 little bit. I think that three years ago ChatGPT
4 didn't exist. Neither did many of the other large
5 foundation models and systems. And now they have, I
6 don't want to be a broken record, but now they have
7 hundreds of millions of users, many of those are
8 children. They're being used and misused in a wide
9 variety of industries and applications because they
10 are general purpose in the way that they're being
11 adopted and used.

12 And so we've already seen a number, if you
13 follow the news, of harms and issues, speculated and
14 real, of how these can have lasting negative effects
15 due to bias, discrimination, all sorts of different
16 things. And so I expect in the next three years we're
17 going to see a lot more of that as adoption becomes
18 wider. There are new applications, and even with the
19 existing applications, people are finding new ways to
20 use or misuse them.

21 There have already been recent research to
22 show that in multilingual uses of these models they're
23 far more vulnerable than in English. And so the
24 multilingual communities that aren't being tested as
25 rigorously as the companies are testing the English

1 prompts are seeing much higher cases of bias toxicity,
2 misinformation, things like that. And so we expect in
3 the next three years that will be many harms that need
4 to be investigated.

5 MR. LI: And please, yes, Dr. Harguess?

6 DR. HARGUESS: Yeah. So I'll be quick. So
7 I think the good news is, is that we're talking about
8 this right now. You know, we haven't had, you know, a
9 huge breach. There are harms that we know about.
10 There are issues with trustworthiness. But, honestly,
11 we haven't really seen the extent of, you know, maybe
12 some of the damage that could be done. So it's an
13 important time to really understand that this research
14 is nascent. We really need to perform this research.
15 And giving these opportunities now versus three years
16 from now is really important.

17 MR. LI: Thank you.

18 And I'll pass this back to my Copyright
19 Office colleagues.

20 MS. KERN: Thank you so much.

21 So we're going to give everybody 15 seconds
22 for closing remarks because we are at 4:30, and then
23 we're going to close out this session.

24 So, if you'd like to give a closing remark,
25 like I said, we'll limit it to 15 seconds and just use

1 the Raise Hand function. And the first person I see
2 is Mr. Reed.

3 MR. REED: Thank you very much for holding
4 this. I want to publicly say that we'd love to work
5 with the proponents to find a solution that would work
6 without having to wade through 1201 procedures every
7 three years, and I would open that up to find some
8 solutions. Thank you.

9 MS. KERN: Mr. Geiger?

10 MR. GEIGER: The good faith AI
11 trustworthiness research is fair use. It contributes
12 to the advancement of computer science. It leads to
13 the production of new creative works. There's no
14 evidence at all that the proposed exemption would
15 result in increased infringement or piracy and, in
16 fact, may well produce the opposite effect by
17 strengthening AI system trustworthiness.

18 Research is being chilled by fear of
19 liability for circumventing TPMs under Section 1201,
20 with more than 350 researchers and journalists calling
21 for protections of this kind of work. The Department
22 of Justice, likewise, said that the security exemption
23 does not likely extend to non-security harms.

24 Our proposed exemption language is carefully
25 crafted to promote public benefit and prevent misuse,

1 prevent infringement modeled closely on existing
2 exemption language for a particular class of users and
3 a particular class of software.

4 I will leave us with a Department of Justice
5 quote. In their letter, they state, "Independent
6 research on the functioning and security of AI systems
7 will likely be essential to ensuring the integrity and
8 safety of AI systems in the future." Thank you.

9 MS. KERN: Thank you.

10 Ms. Cohen? Oh, Ms. Cohen, you're muted.

11 MS. COHEN: Thank you. Independent good
12 faith testing of AI systems for trustworthiness issues
13 is really important to maintain the responsible
14 deployment of AI. And although bias bounty and
15 research access programs can provide AI researchers
16 with permission, those channels are not a replacement
17 for the exemption under Section 1201. They're
18 helpful, but they don't extend to all system
19 providers, nor do they apply to all good faith
20 researchers.

21 So we support the exemption language
22 proposed by the Hacking Policy Council. I think it
23 supports copyright researchers and the public and it
24 uses specific definitions and lessons learned from the
25 highly successful security testing exemption, which

1 has only strengthened our cybersecurity posture.

2 MS. KERN: Thank you.

3 Mr. Longpre?

4 MR. LONGPRE: Yeah. I guess I'd love to
5 leave it on the note about timeliness. I think, while
6 the technology is extremely nascent, the breadth and
7 extent of adoption is so broad and not so nascent.
8 And the effect that it's going to have on people and
9 society is already coming to bear fruit in negative
10 ways, and that compels the importance of good faith,
11 third-party, independent research, which currently we
12 see is being chilled, and we think the evidence on
13 that is fairly clear and this would go a really long
14 way in reducing those chilling effects.

15 MS. KERN: Thank you.

16 Mr. Englund?

17 MR. ENGLUND: I'd just like to underscore
18 that this is a very broad exemption, much broader than
19 what was originally contemplated by the petition or by
20 the NPRM and broader in ways that potentially
21 implicate not only software that is used to secure AI
22 systems but creative content that is also available
23 within those systems and just doesn't seem like the
24 proponents have made a record that would possibly
25 justify the full breadth of the exemption they are

1 seeking.

2 And if the Office applies its traditional,
3 rigorous analysis here, it, I think can't possibly
4 grant an exemption of the breadth that's been
5 requested.

6 MS. KERN: Thank you.

7 Ms. Elazari?

8 DR. ELAZARI: So mitigating the breadth of
9 unintended consequences that can be caused by the wide
10 adoption of AI is one of the highest priorities of
11 this government and this Administration.

12 As Mr. Geiger and my fellow proponents have
13 suggested, we have proposed and brought forward a very
14 careful, with guardrails, with the proper
15 justification type of framework, building on the
16 extensive record of the security research exemption
17 and it's therefore very appropriate to consider this
18 exemption. And the time is now, and we have seen this
19 in other policy action from the Administration.

20 We must equip the security and the auditing,
21 the AI auditing community with the tools that enable
22 them to perform this type of research that we are
23 asking them to do so in other actions in a way that
24 reduces the chilling effect that was well documented
25 on the record.

1 MS. KERN: Thank you. And I will pass it
2 back to the Office's Deputy General Counsel for
3 closing remarks.

4 MS. CHAPUIS: Thanks, everyone, for being
5 here and for your written comments as well. We are
6 adjourned for today. We'll resume 1201 hearings
7 tomorrow at 11 a.m. when we will discuss video game
8 preservation. Thanks.

9 (Whereupon, at 4:36 p.m., the hearing in the
10 above-entitled matter was adjourned.)

11 //

12 //

13 //

14 //

15 //

16 //

17 //

18 //

19 //

20 //

21 //

22 //

23 //

24 //

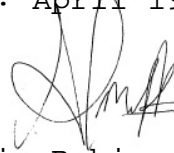
25 //

REPORTER'S CERTIFICATE

CASE TITLE: Section 1201 Public Hearing: Proposed
Class 4, Computer Programs - Generative AI Research
HEARING DATE: April 17, 2024
LOCATION: Washington, D.C.

I hereby certify that the proceedings and
evidence are contained fully and accurately on the
tapes and notes reported by me at the hearing in the
above case before the United States Copyright Office.

Date: April 19, 2024



Alexis Robinson
Official Reporter
Heritage Reporting Corporation
Suite 206
1220 L Street, N.W.
Washington, D.C. 20005-4018