WIKIMEDIA
FOUNDATION

August 19, 2020

The Honorable Thom Tillis, Chairman
The Honorable Christopher A. Coons, Ranking Member
Subcommittee on Intellectual Property
Committee on the Judiciary
United States Senate

Dear Chairman Tillis and Ranking Member Coons:

Thank you for the opportunity to respond to your further questions for the record regarding my testimony before the Subcommittee's July 28 hearing, "How Does the DMCA Contemplate Limitations and Exceptions Like Fair Use?" I have provided my answers to your collected questions below.

Responses to Questions from Chairman Tillis:

1. *I understand that fair use is a judge-made doctrine that Congress codified in section 107 of title 17 while leaving discretion to courts. What are some of the benefits to leaving fair use's development to courts rather than having us define it here in Congress? What are some of the downsides?*

Fair use is famously a fact-specific doctrine, and courts are tailor-made to address and account for the specific facts of a given situation. This traditional role of the common law allows courts to take into account not just broad trends in changes to technology, business practices, and forms of art, education, and commentary,[1] but to account for the specific instances as they arise in factual scenarios, instead of trying to anticipate faster movement of art and culture. Rigid statutory definitions can restrict the progress of art, criticism, and education, especially as educators and critics use more technology that digitally captures our surroundings and creates derivative works.

The open-ended and flexible nature of fair use also is necessary because of the rigid bright lines drawn in Section 106: the fact that legal consequence attaches to such commonplace and ubiquitous activities as public display, public performance, and distribution make it essential for limitations and exceptions like fair use, the first sale doctrine, and various performance exceptions to operate constantly and flexibly. Digital technologies that rely upon copying to operate (every digital device must make

[1] For instance, "supercuts" have become a common form of media criticism, editing together multiple instances of a given trope seen in a particular piece or genre of media. While some supercuts originate from a sense of play or fun, they all can provide the viewer with a unique commentary on how a given trope becomes prevalent in media, and how it fits within the context of the works that use it. The role that a supercut plays in commentary and criticism of media has developed substantially over the past few years.

reproductions of works it is perceiving or processing, even if those reproductions last for mere seconds or are never redistributed as copies), or that capture more of the world around us with ever-present cameras and microphones, mean that fair use must also operate constantly and flexibly to account for the many reproductions that ordinary users make every day, whether that is merely using a web browser or taking snapshots of their street.

The downsides of a lack of statutory specifications are a lack of permanent certainty for those who must make copyright determinations, including stakeholders engaging in Section 512's notice-and-takedown process, where they often must make decisions at scale. Statutory specifications for fair use could make those determinations more definitive, but at a cost. While notice-and-takedown processes would be easier to scale and even automate if fair use could be given a more quantifiable set of criteria, doing so would eliminate the longstanding advantages to fair use's flexibility, including in cases far removed from the notice-and-takedown system, and involving more fundamental rights of freedom of expression. In other words, altering fair use because of problems in Section 512 would be like redefining speech in order to make broadcast decency determinations easier to automate.

Businesses naturally want a well-defined set of instructions in the law that they can follow like an algorithm. That makes solutions easy to scale, whether by automation, or providing a simple checklist for a new hire to follow when screening uploads. However, this naturally creates regular problems when reality does not fit the predetermined assumptions of those who set up the algorithm or the checklist. This creates a sort of "sorcerer's apprentice" problem[2] where the bureaucratic system, unable to account for changing circumstances, continues to operate, leading to absurd and frustrating results. This can be the case with a runaway algorithm, but also takes place in human systems governed by overly rigid rules, as experienced by anyone trying to call a customer service line that is staffed by a human, but one who is required to follow an excessively constrained script.

However, this does not mean that fair use determinations must necessarily be murky or only made by legal experts: the broader contours of fair use are largely undisputed, outlined by key appellate and Supreme Court cases, and can be relied upon without additional statutory specification. The statute itself also contains specific examples of clear fair use scenarios beyond the generalizations of the four enumerated factors that can be applied quickly even if they resist certain types of automation.

2. *In Lenz v. Universal Music Corp., the Ninth Circuit held that a copyright owner must affirmatively consider the existence of a fair use defense before sending a takedown notice or else risk liability under section 512(f). What does it mean to consider fair use? What are some practices that copyright owners are using to do that? I'm concerned with making sure that section 512 scales in today's digital world. Under Lenz, can copyright owners still do automated notice-sending to keep up with the volume of infringement?*

[2] *See* https://en.wikipedia.org/wiki/The_Sorcerer%27s_Apprentice.

To "consider fair use" is a relatively low requirement. Notably, this is not the same as conducting the sort of analysis that a district court would in making a finding of fair use. Instead, it is making sure (among other things) that a notice sender is not sending notices for obvious fair uses.

Because fair use is multifaceted, there are ways in which notice senders are already able to account for certain factors that are highly relevant to fair use, even if few, if any, are able to conduct a full analysis. For example, if a notice sender is telling its employees or its algorithms to not include cases where the work is being played or displayed in the background of a candid video shot, it is already considering a type of fair use. If the notice sender automatically excludes matches of less than a few seconds, that is also approximating some aspects of fair use. Considering fair use requires more than those minimum standards, but they are a necessary part of it.

The legal assistant who performed the initial analysis for Universal in *Lenz* was working from criteria that approximate some aspects of fair use, but also clearly missed other, major aspects of it, including purpose and character. As the opinion notes, the assistant was to determine if the song was "recognizable" and "the focus of the video." These are reasonable, if demonstrably insufficient, factors to consider in a full fair use determination. Under the Ninth Circuit's subjective good faith belief standard, criteria like those that Universal employed would be considered a sufficient consideration of fair use.

Despite this being the legal standard set in the Ninth Circuit, I believe that notice senders can develop better-tailored criteria for fair use than those used by Universal several years ago in *Lenz*. Some of these criteria may be built into automated systems; others can be scaled by augmenting human work via automation.

By way of analogy, Wikimedia Foundation staff and volunteers both create automated tools that assist volunteers in editing Wikipedia. However, where the tools are meant to account for complex, fact-specific determinations (like article quality), instead of simple, bright-line determinations (like identifying broken links), those tools do not make direct edits, but rather flag likely issues for human review. The automation serves as a way to cull the to-do list of a human editor exercising context-specific judgment, instead of presuming that the algorithm will understand the context. This requires not just human review, but also that the human doing the review has a basic understanding of fair use and can sort the easy cases from the difficult ones, allowing notices to be sent or not for clear-cut cases, and for a more exacting review to be applied to closer calls.

3. *Is the counter-notification process sufficient to protect fair use's role in section 512? Are service providers able to counsel users or encourage them to file a counter-notice? If so, do you know how frequently they ever do that?*

The counter-notification process is not sufficient to protect fair use. As the Copyright Office, Professor Ginsburg, and others have frequently noted, counter-notices are rare, and especially rare in comparison to the volume of notices. One reason for this would seem to be that the benefits of fair use are very frequently widely distributed among the public generally. While a commercial parody might benefit the parodist more specifically, fair use is much more commonly benefitting reporting, education,

commentary or criticism that is not directly benefiting any one person, but rather society as a whole. That distributed benefit is not likely to outweigh the risk of litigation faced by any one potential counter-notice sender.

If the counter-notification process were sufficient to protect fair use, the Foundation would not feel the need to conduct our own fair use determinations when we receive a takedown notice, which both requires more staff time than processing a counter notice and also exposes the Foundation to liability risk that would be absent if we simply followed the bare requirements of Section 512.

Despite the fact that the Foundation chooses to do this does not mean that we recommend others do the same. Nor would we want OSPs to counsel users that they should file counter-notices. The limitations on liability for OSPs in Section 512 are designed with the premise that OSPs will act to minimize their copyright infringement liability, particularly given the high costs of litigation and statutory damages. An OSP advising a user to file a counter-notice would be encouraging that user to take on litigation risk that the OSP, in many cases, would itself be hesitant to bear. This could raise several potential conflict of interest and other ethical issues, and would seem to require a much closer relationship between the user and the OSP.

Making determinations of whether a particular use of a work was authorized, either by a rightsholder or by the law, requires facts usually held not by an OSP, but either by the rightsholder or the uploader. Further, making a determination of whether a an uploader should file a counter-notice would require facts that will largely be unknown to the platform, including knowledge of the uploader's confidence in their purpose and character of the use; their risk tolerance for potential liability, and ability to fairly litigate their case.

4. *Section 512(m) says that service providers have no duty to monitor for infringement – does that play into whether service providers are willing to actively evaluate alleged infringements for fair use?*

Section 512(m) does not play a direct or large role in evaluating alleged infringements for fair use. With or without Section 512(m), an OSP lacks an incentive to question the judgment of a notice or counter-notice sender. A hypothetical affirmative duty to monitor for infringement would not seem to incentivize an OSP to question the results of its monitoring any more than it would question notices under the current statute.

5. *You mentioned that the Wikimedia Foundation directly reviews takedown notices to see if the alleged infringement might be a fair use. You also said this was possible because the Foundation and Wikipedia receive so few takedown notices. What do you think the role of OSPs should be in responding to takedown notices that address works that could be fair use?*

The Wikimedia Foundation seems to differ from the average OSP in two respects: first, the small volume of takedown notices that we receive due to the focused mission of our projects and the diligence of the volunteer community, and second, the mission-driven impulse to ensure that free knowledge is

widely accessible, including under fair use. Our role in responding to takedown notices that could be fair use is therefore similarly distinct from most OSPs', even though we take that role seriously.

Section 512 is designed to avoid requiring OSPs to make infringement determinations (including fair use determinations) as a condition of limitations on liability; if OSP liability were premised on adjudicating the merits of these cases, very few would be resolved in a timeframe satisfactory to any party. As the case law on fair use shows, even district and appellate courts are reversed on complex cases of fair use; the courts themselves do not suffer liability or penalties upon being reversed by higher courts.

6. *As Professor Ginsburg noted in her testimony, fair use is not a defense to an act of circumvention in violation of section 1201. This was a conscious decision that Congress made in 1998, instead enacting some permanent exemptions and a rulemaking process to establish temporary exemptions every three years. Do you think that the rulemaking process is doing the necessary work for fair use?*

The triennial process is currently a useful step that helps accommodate fair use within the anti-circumvention provisions, but those provisions remain far more restrictive than copyright law itself. The exemptions created in the triennial process apply only to the circumvention of access controls under Section 1201(a)(1); there is no similar mechanism for anyone to develop exemptions for the creation or distribution of access circumvention tools under Section 1201(a)(2) or for the creation or distribution of copy-protection circumvention tools under Section 1201(b). That oversight seems based in part upon the mistaken premise that the only people who would need to circumvent these controls would be well-financed and technologically sophisticated actors. As seen in many of the past triennial processes, this is not necessarily the case, as media students, disabled people, and others, of widely varying technological, financial, and legal resources have need of these tools. The exclusion of sections 1201(a)(1) and 1201(b) from the scope of the triennials eliminates an open market for such tools.

Beyond this, the triennial process itself is an imperfect mechanism for accommodating fair use, which does not require any administrative process to operate. Even as we observe the difficulty in ensuring that non-copyright experts have in understanding the various statutes at issue, understanding the rules and import of any administrative process is even less well-known and less accessible to the average user. While the Copyright Office, Library of Congress, and NTIA have made improvements to the process over the past several rounds of rulemaking (including, notably, the renewal of past-granted exemptions in the absence of opposition), the process still remains far more cumbersome than a system that recognized that circumvention for non-infringing purposes was itself not a violation.

Responses to Questions from Ranking Member Coons:

1. *You testified that, at least in some cases, the fair use analysis can be too complex for automated tools. Could automated technological measures nonetheless serve to filter the worst digital piracy offenders and alleviate the burden of processing large numbers of takedown notices while leaving the tricky cases to humans?*

Automated processes are certainly being used today in many contexts, both by notice senders and by OSPs, to speed the notice-and-takedown process. The difficulty is not in the existence or use of these tools, but how they are deployed and what is done with the results that they generate. Where automated tools are used to sort through large volumes of data to highlight likely candidates for human review, they can speed the process without introducing additional bulk errors. However, different tools are designed for different contexts, and with different platforms as their model. Just as a manifesto calling for rebellion has a different purpose and character when published in a broadsheet versus displayed in a museum or in a history book, the upload of a given file can have a different purpose when uploaded to a general-purpose sharing platform versus an educational archive. Human judgment goes a long way to incorporating those types of contextual cues into infringement determinations.

As noted in my response to Chairman Tillis's Question 2 above, automated tools can be useful, in our experience, in helping humans deal with even complex determinations and make them more efficiently. This does, however, require a strong fit between the tool and the context of its use, and a good understanding of both by the human deploying the tool to be both effective and efficient. Mandating the use of automated tools can force their use in inapposite contexts, further increasing the potential for error.

2. *Fair use is a particularly thorny legal doctrine. Is it reasonable to expect typical creators and internet users to understand and apply the fair use factors as part of the notice-and-takedown process?*

Many creators, rightsholders, and internet users already do account for fair use in a number of ways, though mostly only partially. While it is unrealistic to expect a rightsholder to make a fair use determination with the same rate of accuracy or level of diligence as a district court, there are clear cases of fair use that should not generate takedown notices. Bad-faith and fraudulent notices occupy one end of the spectrum, but other clear cases can include instances where any reasonable human review would have revealed that the copyrighted work in question was not being used (such as when the same public domain work is used in the copyrighted work and in the alleged infringement), or that the use was clearly nonsubstantial and incidental. Notice senders can and should also exercise more caution when there appears to be news reporting, commentary, or criticism at stake; the fact that these may sometimes be harder calls is counterbalanced by the fact that these are often the types of uses that can most harm the public interest if they are mistakenly taken down.

Making fair use determinations in an informal, extrajudicial context need not require the exacting legal analysis that a court is required to. As you note in Question 1, there are easy cases as well as hard ones; easy cases for fair use can be readily identified and removed from a notice-sending queue based upon rubrics that track Section 107's more specific examples of "criticism, comment, news reporting, teaching (including multiple copies for classroom use), scholarship, or research." Such guidelines can be implemented efficiently by a notice sender, even if they will necessarily be less precise than a formal adjudication. Ideally, notice senders should not be penalized for getting a thorny doctrine wrong, but they also should not be free to ignore the doctrine in cases where it obviously applies. More careful evaluations that consider fair use can prevent mistaken notices sent for the large number of incidental fair uses, and

can also prevent the rarer, but more pernicious problem of notices sent for fair uses that are news reporting, criticism, commentary, or parody.

Unlike notice senders, uploaders whose files have been taken down already have a strong incentive to clearly understand the contours of fair use before sending a counter-notice; since a counter-notice is often essentially the next step in progress to litigation, they have a specific reason to ensure that their actions are authorized under the law before proceeding.

a. What liability – if any – should they face if they get it wrong?

Section 512(f) provides an existing framework for liability; its text does require knowledge that the notice was materially a misrepresentation, ensuring that reasonable, good-faith notices should never result in penalties. The standard currently set by the Ninth Circuit in *Lenz* is even more forgiving, applying a subjective standard of good faith, meaning that even an unreasonable belief that the notice was not a misrepresentation could result in no penalty. An objective standard can help ensure that those parties best placed to evaluate the merits of their notices take reasonable steps before alleging infringement that will result in a takedown.

b. Are online service providers better situated to evaluate fair use?

Online service providers are usually not as well situated to evaluate fair use. An OSP has less information than the rightsholder about the identity or metes and bounds of the original work (for instance, if the notice sender is claiming a copyright on an underlying musical work, or on the particular sound recording). The rightsholder also will have more information about the characteristics of the original work, such as its length, its main themes, and thus the proportion of those used in the upload. In other words, the rightsholder has more information about the nature of the original work, and more information about the amount and substantiality of the use. That information is critical in evaluating the purpose and character of the use and the effect of the use on the potential market for the original work, giving the notice sender substantially more information about all four factors than the OSP.

3. To what extent does the risk of losing safe harbor protection deter online platforms from performing their own fair use analysis when evaluating takedown notices or counter notices?

Though I cannot speak for the motivations of other platforms, I believe that this is substantial. OSPs follow the strictures of Section 512 in order to ensure that they are not liable for the potentially massive statutory damages that could result from a multitude of works being infringed by their users. Conducting a fair use analysis, and then using the results of that analysis to refuse to comply with a takedown notice, creates a completely avoidable litigation risk for the OSP.

OSPs would also be exposing themselves to litigation risk by refusing to comply with a counter-notice, although the potential liability is often lower or less definite, since it is less likely that the counter-notice sender would be able to allege copyright infringement by the OSP.

4. *The Copyright Office has recommended that we reject a one-size-fits-all approach to modern internet policy. How should the differences among stakeholders influence our evaluation of fair use in the context of the DMCA?*

Flexibility and adaptability are already built into the contours of fair use; each of the four required factors is meant to allow a court to account for different types of uses, works, contexts, and market effects. The paradigmatic examples of fair use specified in Section 107 also naturally incorporate mechanisms to account for different types of uses and settings, such as specifically reproductions for use in classrooms.

More specific definitions of different types of stakeholders can in many cases create better fits between laws and desired policy outcomes, but the level of specificity best suited for a statute is highly determined by its scope and nature.

5. *I hear stories from rights holders who file millions of takedown notices every year. To what extent should we expect them to perform a fair use analysis for each such notice?*

It can be reasonable to expect a certain amount of fair use analysis for every such notice. That doesn't necessarily mean the sort of certainty that would be required of a court making a determination that would withstand appeal; but it can include a type of analysis that should be enough for obvious cases. We can expect a notice sender to not send notices when the use is incidental background inclusion; when an algorithm was triggered by an audio or video clip not owned by the notice sender (such as a public domain work used by both works at issue); or when the use is a clear case of parody, criticism, or commentary. As *Lenz* has shown, at least some consideration of fair use is required by the statute as written; it makes sense that notices should not be sent when a reasonable person would believe that the use was fair.

* * * *

Once again, thank you for the opportunity to testify at the hearing and append my responses to your questions here. If I can provide you with any more information, please do not hesitate to contact me.

Best regards,

Sherwin Siy
Lead Public Policy Manager
Wikimedia Foundation