

FOR PUBLICATION

In the
United States Court of Appeals
For the Eleventh Circuit

No. 23-13156

ATHOS OVERSEAS LIMITED CORP.,

Plaintiff-Appellant,

versus

YOUTUBE, INC.,
YOUTUBE, LLC.,
GOOGLE, LLC.,

Defendants-Appellees.

Appeal from the United States District Court
for the Southern District of Florida
D.C. Docket No. 1:21-cv-21698-DPG

Before WILLIAM PRYOR, Chief Judge, and JORDAN and MARCUS, Circuit Judges.

JORDAN, Circuit Judge:

This case concerns the application of one of the safe-harbor provisions of the 1998 Digital Millennium Copyright Act, 17 U.S.C. § 512(c), in the modern digital age.

Athos Overseas Limited, which owns the copyright to many classic Mexican and Latin American films, sued YouTube, Inc., YouTube, LLC, and YouTube's owner, Google, LLC (collectively "YouTube"), alleging copyright infringement based on the unauthorized posting of its copyrighted material on the YouTube website, an internet platform where users can upload video content for public viewing by visitors of the site. Following discovery, Athos filed a motion for partial summary judgment, and YouTube filed its own motion for summary judgment. The magistrate judge issued a report recommending that Athos' motion be denied and that YouTube's motion be granted based on § 512(c), a safe-harbor provision of the DMCA. The district court adopted the report and recommendation and entered final judgment in favor of YouTube. This appeal followed.

After review of the record and the parties' briefs, and with the benefit of oral argument, we affirm. We agree with the district court that on this record YouTube was protected by § 512(c).

23-13156

Opinion of the Court

3

I

In 1998 Congress, foreseeing the significance the internet would play in the coming millennium, enacted the DMCA “to update domestic copyright law for the digital age.” *Viacom Int’l Inc. v. YouTube, Inc.*, 676 F.3d 19, 26 (2d Cir. 2012). The DMCA was meant to “provide certainty for copyright owners and Internet service providers with respect to copyright infringement liability” by creating protections from copyright liability under certain circumstances and by establishing a regime through which copyright owners could reliably enforce their rights. *See* S. Rep. No. 105-190, at 2 (1998). The DMCA sought to balance the interests of copyright owners and those of the nation by ensuring “that the efficiency of the Internet [would] continue to improve and that the variety and quality of services on the Internet [would] expand.” *Id.* *See also* 4 Melville B. Nimmer & David Nimmer, *Nimmer on Copyright* § 12B.01[C][1] at 12B-28 (2020) (“To balance these divergent interests, the [DMCA] as enacted embodies disparate forms of protection.”).

Congress achieved this balance in part by establishing specific protections for internet service providers. The DMCA safe-harbor provision at issue here applies to “information residing on systems or networks at direction of users” and reads as follows:

(1) In general.—A service provider shall not be liable for monetary relief, or, except as provided in subsection (j), for injunctive or other equitable relief, for infringement of copyright by reason of the storage at the direction of a user of material that resides on a

system or network controlled or operated by or for the service provider, if the service provider—

(A)(i) does not have actual knowledge that the material or an activity using the material on the system or network is infringing;

(ii) in the absence of such actual knowledge, is not aware of facts or circumstances from which infringing activity is apparent; or

(iii) upon obtaining such knowledge or awareness, acts expeditiously to remove, or disable access to, the material;

(B) does not receive a financial benefit directly attributable to the infringing activity, in a case in which the service provider has the right and ability to control such activity; and

(C) upon notification of claimed infringement as described in paragraph (3), responds expeditiously to remove, or disable access to, the material that is claimed to be infringing or to be the subject of infringing activity.

17 U.S.C. § 512(c)(1). Significantly, the safe-harbor provisions of the DMCA, including § 512(c), are not conditioned on “a service provider monitoring its service or affirmatively seeking facts indicating infringing activity[.]” 17 U.S.C. § 512(m)(1).

Through § 512(c)(1)(C), Congress also established a “notice and takedown” regime for the benefit of copyright owners on the other side of the scales. *See* S. Rep. No. 105-190, at 45. To be effective, a copyright owner’s takedown request must include

23-13156

Opinion of the Court

5

“[i]dentification of the material that is claimed to be infringing . . . and information reasonably sufficient to permit the service provider to locate the material.” § 512(c)(3)(A)(iii). See *Lenz v. Universal Music Corp.*, 815 F.3d 1145, 1151 (9th Cir. 2016) (explaining that a takedown notification must include “identification of the copyrighted work, identification of the allegedly infringing material, and, critically, a statement that the copyright holder believes in good faith the infringing material is ‘not authorized by the copyright owner, its agents, or the law’”); *Perfect 10, Inc. v. CCBill LLC*, 488 F.3d 1102, 1112 (9th Cir. 2007) (same).¹

Under this notice-and-takedown regime service providers are at risk of losing the protection of the safe-harbor provision should they fail to adequately respond to a claim of copyright infringement. See § 512(c)(1)(C). Service providers may also lose safe-harbor protection if they neglect to take down infringing material they have actual or red flag knowledge of, regardless of whether any takedown request has been submitted. See § 512(c)(1)(A)(i)-(iii).

Congress intended the notice-and-takedown regime as “a formalization and refinement of a cooperative process that ha[d] been employed to deal efficiently with network-based copyright infringement.” S. Rep. No. 105-190, at 45. It has been described as “a

¹ Some copyright owners have complained that the DMCA’s notice-and-takedown regime “more closely resemble[s] the game of Whack-A-Mole than an efficacious tool for relief.” 4 Nimmer on Copyright § 12B.01[C][5] at 12B-32.1. Needless to say, we apply § 512(c) as it was written, not as it could have been written.

‘compromise’ between protecting copyright owners and ‘insulat[ing] service providers from liability for infringements of which they are unaware . . . so as to make it commercially feasible for them to provide valuable Internet services to the public.’” *Capitol Recs., LLC v. Vimeo, Inc.*, 125 F.4th 409, 413 (2d Cir. 2025) (*Vimeo II*) (citing and quoting *Capitol Recs., LLC v. Vimeo, LLC*, 826 F.3d 78, 82 (2d Cir. 2016) (*Vimeo I*)).

To qualify for safe-harbor protection under § 512(c), a service provider like YouTube (1) must be a covered service provider under § 512(k)(1); (2) must adopt and reasonably implement (and inform subscribers and account holders of) a policy providing for termination of those who are repeat infringers as set out in § 512(i)(1)(A); (3) must expeditiously remove any infringing material it becomes aware of on its website, whether through its own actual or red flag knowledge or through receipt of a valid takedown notice; and (4) must not directly financially benefit from infringing material on the website if it has the right and ability to control that material. See § 512(c)(1); *Vimeo II*, 125 F.4th at 413–14; *Perfect 10*, 488 F.3d at 1109.

The safe-harbor set out in § 512(c) is “properly seen as an affirmative defense.” *Vimeo I*, 826 F.3d at 94. Athos acknowledges that YouTube is a service provider generally covered by the § 512(c) safe-harbor provision. And it does not claim that YouTube does not expeditiously remove the specific content for which a copyright owner sends an effective DMCA takedown request. Instead, Athos attacks YouTube’s protection under § 512(c) on two fronts.

First, Athos maintains that YouTube has actual or red flag knowledge of additional infringing material on its website or is otherwise being willfully blind to such material, and that material is not being expeditiously removed. Second, Athos asserts that YouTube derives a direct financial benefit from infringing material posted to its website while also having the right and ability to control that material. Athos' arguments are in part based on certain factual assertions relating to particular technologies and functionalities of the YouTube website. Athos believes that these technologies and functionalities deliver sufficient knowledge of infringing material to YouTube and facilitate an ability to control that material on the site such that YouTube should lose its § 512(c) safe-harbor protection.

II

We review a district court's grant of summary judgment *de novo*, viewing the facts and making inferences in favor of the non-moving party. *See Smith v. Owens*, 848 F.3d 975, 978 (11th Cir. 2017). Summary judgment under Rule 56 is appropriate where no material facts are in genuine dispute and one party is entitled to judgment as a matter of law. *See id.* An issue of fact is material if it has the potential to affect the resolution of the case under the applicable legal principles. *See Harrison v. Culliver*, 746 F.3d 1288, 1297–98 (11th Cir. 2014). And a fact is genuinely in dispute if, taking the record as a whole, a rational trier of fact could find in favor of the nonmoving party. *See id.* at 1298.

Athos challenges only the district court’s grant of YouTube’s motion for summary judgment and not the denial of its own motion for partial summary judgment. As a result, we set out the facts, and draw all reasonable inferences, in favor of Athos. Yet, as we explain, some of Athos’ factual assertions are not supported by the record.

III

We start by describing the specific technologies and functionalities that are at issue here—what they are, what they do and don’t do, and how they do it—before we wade into application of § 512(c) of the DMCA. *See, e.g., Vimeo II*, 125 F.4th at 416 (explaining Vimeo’s curation and moderation functionalities before evaluating issues relating to equivalent red flag knowledge and right and ability to control). That description is critical because the parties characterize the relevant technologies and functionalities differently in their briefs.

Two suites of technologies and functionalities are at issue here. The first suite encompasses the functionalities surrounding YouTube’s moderation and curation of the content on the YouTube website, which are relevant to whether YouTube has the right and ability to control infringing material. The second suite is made up of YouTube’s video-hashing technology as well as the multiple video-hash-matching technologies housed within its various copyright management tools.

23-13156

Opinion of the Court

9

A

The first suite of functionalities relates to the ability of YouTube to moderate and curate user uploads to its website.

Athos contends that YouTube exerts substantial amounts of “control over the users who upload the infringing material, its users in general and the infringing activity on its platform.” Appellant’s Br. at 48 (citing D.E. 137-3 at 15:1–7, 17:12–25, 79:24–87:25, 119:18–24, 120:24–121:04; D.E. 137-6 at 102:25–103:10, 118:17–120:09, 132:08–134:21, 191:11–12; D.E. 137-7 at 53:08–57:20). After careful review of Athos’ record citations, there are only two that reference YouTube’s ability to control the content on the YouTube website. *See* D.E. 137-6 at 132:08–134:21 (explaining how a user can use the Prevent Copies functionality of YouTube’s copyright web form to have its system prevent the reupload of exact and near-exact copies of already removed videos); D.E. 137-7 at 53:08–57:03 (discussing YouTube’s ability to suggest videos to users and to autoplay videos after one video has finished, and confirming that YouTube can take down any page on its website as well as set its own policies and procedures for what content can be uploaded and viewed).

These matters are not in dispute. YouTube does not contest that its Prevent Copies functionality exists, or that its algorithmic video curation can suggest and autoplay other videos available on the site, or that it has the general power to remove content from the site.

B

YouTube currently uses video-hashing technology to create an identifying chain of characters, i.e., a hash, for each video uploaded to its website. The video hashes produced by YouTube's technology are used to facilitate video-hash-matching functionalities available through various copyright management tools that YouTube offers to users. *See* D.E. 137-6 at 144:16–145:08 (explaining that there are three different copyright management tools offered by YouTube to users that rely on some form of hash-matching technology: (1) Content ID, (2) the Copyright Match Tool, and (3) the Prevent Copies functionality of its copyright web form).

1

The various hash-matching functionalities are each housed in different copyright management tools, but they all have the common capability of comparing a reference video hash to YouTube's video-hash database to produce a list of video-hash matches that can then be processed in some other manner. The primary differences between the tools for our purposes are the options available to the user for determining what, if any, next steps should be taken with respect to the list produced by the video-hash-matching process. *See* D.E. 14-2 (discussing Content ID, which allows the user to select from various options, including automatic removal, various monetization settings, or tracking of flagged matches); D.E. 136-3 at 149:14–150:07 (explaining the Prevent Copies functionality, which can prevent the subsequent upload of exact or near-exact copies of already removed videos if a user checks the applicable

box on a submitted copyright web form); D.E. 136-1 at 18–19 (explaining the Copyright Match tool, which only presents the user with the identified matches for the user to then submit takedown requests for those matches the user believes to be infringing after its own review). Only the Prevent Copies functionality is available to all users; the other tools’ availability is based on a user’s “[d]emonstrated need,” “[a]vailable resources to manage [its] rights and content,” and “[k]nowledge of YouTube’s copyright system.” D.E. 14-3 at 2. This is due to YouTube’s concerns for “protecting against significant disruptions that can result from [these functionalities’] misuse.” *Id.*

YouTube offers its copyright management tools for use by users and does not operate any of them automatically itself because it is the users, i.e., the copyright owners, who have the best knowledge of, and access to, their copyright information. The users must decide what each tool should do with an identified match, as it is their underlying copyrights that are being protected. In the case of Content ID, the user has broad and precise powers to direct the takedown, tracking, or monetization of identified video-hash matches on a region-by-region basis. Without input from a user regarding which actions to take in which regions, YouTube has no way of administering Content ID on the user’s behalf.

Notably, with respect to YouTube’s hash-matching technology, it is possible that a video-hash match identified by the hash-matching process is not actually a violation of the underlying copyright. This can happen, for example, when the use of the

copyrighted content generating the match constitutes fair use or when it is uploaded by the copyright owner, a subsidiary, or a licensee. *See* D.E. 137-6 at 74:24–25 (testimony of Kevin Zhu, a Content ID product manager: “I don’t think the system has the capacity to determine potential infringement.”).

There is no evidence in the record that YouTube has any tool capable of performing a legal evaluation to determine whether an identified video-hash match is actually infringing on a given copyright. And the record indicates that YouTube seeks to avoid mass removals of legal uses of copyrighted material by its users. As a result, none of YouTube’s copyright management tools operate automatically without substantial input and management by the copyright owners, and the most powerful tools are only available to qualifying users.

YouTube offered Athos the use of its various copyright management tools, including Content ID, but Athos declined the offer because it viewed the required user agreement as “monopolistic and abusive.” Athos chose instead to hire a law firm to manually locate potential instances of infringement and send individualized takedown requests to YouTube. Neither Athos nor its counsel ever used Content ID or the Copyright Match tool to assist in the management of its copyrights on the YouTube website.

2

According to Athos, the second suite of technologies is made up of various functionalities that all fall under the umbrella term Content ID. Athos contends this is a system on the YouTube

website that “automatically generates a digital fingerprint of every video” (what Athos refers to as digital or video fingerprints in its briefing we refer to as “video hashes”) at the moment it is uploaded to the YouTube website. *See* Appellant’s Br. at 9. Athos asserts that this system “mechanically and instantaneously compares the fingerprint with other fingerprints in the platform for purposes of obtaining matches of material that copyright owners have informed YouTube is copyrighted.” *Id.* Athos maintains that these functionalities are run automatically on all videos uploaded to the YouTube website. *See id.* at 9–10. As a result, Athos believes that at the moment YouTube receives a takedown request premised on Athos’ ownership of a given work appearing at a particular location on the YouTube website, YouTube’s technologies immediately provide YouTube actual or red flag knowledge of additional infringing material in the form of any locatable video-hash matches to the identified infringing material. *See* Oral Argument Audio at 12:00–13:47.

For this theory, Athos relies exclusively on its interpretation of the deposition testimony of YouTube’s representatives. But the record taken as a whole (and viewed in Athos’ favor) does not support Athos’ characterization of YouTube’s copyright management tools. Our review of the record shows that Athos’ factual assertions rely primarily on unsustainable readings of the deposition of Mr. Zhu, a Content ID product manager.²

² We recognize that Mr. Zhu’s deposition, as well as certain portions of Athos’ brief, were sealed pursuant to the parties’ stipulated protective order. *See* D.E.

As an example, in its initial brief Athos says that “[r]emoving all existing infringing content as opposed to removing clip by clip, whenever a copyright owner identifies a specific URL, could ‘shut down the whole site.’” Appellant’s Br. at 11 (citing and quoting D.E. 137-6 at 200:10–11). When asked about the technical feasibility of taking down all near-exact hash matches of videos identified by a URL in a DMCA takedown request, Mr. Zhu stated: “I’m not even sure if it’s ultimately feasible to do in a way that makes sense. So I think it’s a little difficult to answer the question because nearly anything is possible with software. Like, we could shut down the whole site, you know, potentially, but that doesn’t mean that would be a good thing to do.” D.E. 137-6 at 199:14–200:12. Mr. Zhu’s observation offers no support for the factual proposition Athos attempts to bolster with it—that YouTube refuses to remove infringing material for fear of shutting down the website.

Athos’ other record citations do not allow a reasonable jury to find that “YouTube automatically compares each [video] fingerprint to all content in YouTube, as a business policy.” Appellant’s Br. at 15 (citing D.E. 137-6 at 103:7–10, 118:17–120:08). For instance, in response to questions regarding whether any aspects of the video upload process on the YouTube website involve human interaction by YouTube employees, Mr. Zhu stated: “I don’t think I know—so I’m not personally specifically aware of processes where a human interacts with a video during that process, but that

60. We quote both carefully to avoid the disclosure of confidential and proprietary information.

doesn't necessarily mean it doesn't exist." D.E. 137-6 at 120:03–08. A fuller review of the second record citation and surrounding context reveals that, after first clarifying that "there are some conditions that might determine when exactly would [Content ID] look for those [hash-matched] videos after the reference file was provided," and that YouTube requires Content ID partners to upload reference files themselves for Content ID to operate, Mr. Zhu explained that "[w]hen a video is uploaded . . . a fingerprint is generated, and a fingerprint is also generated by Content ID reference file, and those fingerprints are compared to each other, as I understand it." D.E. 137-6 at 100:10–103:10. At no point did Mr. Zhu state that Content ID operates without a reference file being provided by the Content ID user. Nor did he suggest that YouTube operates Content ID universally or automatically as a business policy without any prior input from a Content ID user.

The record citations that Athos relies on also do not support the factual proposition that "[t]he purpose of comparing fingerprints is ensuring YouTube locates and flags infringing content." Appellant's Br. at 16 (citing D.E. 137-6 at 70:17–20, 146:12–147:03). Mr. Zhu's first cited deposition excerpt directly contradicts the proposition Athos asserts. *See* D.E. 137-6 at 70:17–20 ("The purpose of the [hash-matching] tool is to show videos that are potential matches and allow the user to decide whether they want to request removal of those videos or do something else."). The second deposition excerpt does nothing to shore up Athos' unfounded assertion; when Mr. Zhu was asked whether he knows if the different hash matching copyright management tools YouTube has made

available to users all rely on the same hashing technology, he gave this answer: “I don’t personally know if there are differences or what differences there might be . . . I’m not the technical expert in the matching technology itself, so it doesn’t necessarily mean there aren’t differences.” D.E. 137-6 at 146:12–147:03.

3

In sum, as to curation and moderation technologies, the record reflects agreement between the parties that YouTube can algorithmically suggest and auto-play videos, can remove users and videos from the site, and can set its own guidelines for users and uploaded content on the site.

With respect to the hashing technologies, the record does not support Athos’ factual assertions. YouTube does not run Content ID (or any of its other copyright management tools) automatically or universally for all videos uploaded to the YouTube website. Athos, moreover, has offered no evidence that any of YouTube’s video-hash-matching or other tools have any functionality to perform a legal analysis of whether an identified video-hash match constitutes copyright infringement. In fact, YouTube has presented un rebutted evidence that all of its copyright management tools were designed with user input and management in mind because it would be unfeasible for YouTube to operate the tools by itself without such input. Athos simply has not identified any material facts in dispute.

III

According to Athos, once it notifies YouTube of its ownership of a given copyright and identifies a single location (e.g., a URL) where the copyrighted material appears on the YouTube website, YouTube then has a video hash sufficient for it to locate additional infringing material on the site. Athos argues that this confluence of technologies generates actual or red flag knowledge of infringement that should result in YouTube being required to remove any match generated by the hash-matching technologies or otherwise risk losing its safe-harbor protection under § 512(c) of the DMCA. Athos is in effect demanding that YouTube be required to administer the operation of its copyright management tools in a way that it never has before and in a manner that it did not design them for, on behalf of and without any input from a user other than a single DMCA takedown request.

Athos challenges the district court’s reliance on Second and Ninth Circuit cases considering equivalent issues—*Viacom Int’l, Inc. v. YouTube, Inc.*, 676 F.3d 19 (2d Cir. 2012), and *UMG Recordings, Inc. v. Shelter Cap. Partners LLC*, 718 F.3d 1006 (9th Cir. 2013)—because it contends that the approach in these cases “disproportionately protects service providers, disincentivizes the creation of creative work and diminishes the value of obtaining copyright protection.” Appellant’s Br. at 27. We are not persuaded.

A

The safe-harbor provision set out in § 512(c) encompasses both an actual knowledge standard and what has become known

as a red flag knowledge standard. *See* § 512(c)(1)(A)(i)–(ii). “[T]he actual knowledge provision turns on whether the provider actually or ‘subjectively’ knew of specific infringement[.]” *Viacom*, 676 F.3d at 31. The “red flag provision turns on whether the provider was subjectively aware of facts that would have made the specific infringement ‘objectively’ obvious to a reasonable person.” *Id.*

The Second Circuit in *Viacom* concluded that the removal obligation in § 512(c)(1)(A)(ii) requires more than a vague knowledge and awareness by a service provider of the general existence of infringing material on its website for the provider to lose safe-harbor protection—the provider must also have not acted “expeditiously to remove, or disable access to, the material[.]” § 512(c)(1)(C). The Second Circuit reasoned that the basic operation of § 512(c) “contemplates knowledge or awareness of specific infringing material, because expeditious removal is possible only if the service provider knows with particularity which items to remove.” 676 F.3d at 30. The specificity requirement therefore applies with equal force to red flag knowledge under the DMCA. *See id.* at 31. As a result, any alleged actual or red flag knowledge would need to relate to specific instances of infringement to trigger YouTube’s obligation to expeditiously remove the content.

The Ninth Circuit came to the same conclusion in *UMG*, 718 F.3d at 1022. It held “that merely hosting a category of copyrightable content, such as music videos, with the general knowledge that one’s services could be used to share infringing material, is insufficient to meet” either the actual or red flag knowledge

standards in § 512(c)(1)(A). *See id.* at 1022–23. The service provider in that case did not lose safe-harbor protection because it “promptly removed infringing material when it became aware of specific instances of infringement.” *Id.* at 1023.

1

Under *Viacom* and *UMG*, YouTube is entitled to safe-harbor protection under § 512(c) because Athos agrees that it expeditiously removes infringing material identified in a takedown request. But this approach to the specificity standard applicable to red flag knowledge has been questioned by a leading copyright treatise. *See* 4 Nimmer on Copyright § 12B.04[A][1][b][ii] at 12B-58–59 (asserting that *Viacom* and *UMG* present an incorrect approach to red flag knowledge). According to Nimmer, “the ‘actual knowledge’ prong is reasonably construed to refer to *specifics*, whereas the ‘red flag’ prong deals with *generalities*.” *Id.* Nimmer asserts that “to show that a ‘red flag’ disqualifies defendant from the safe harbor, the copyright owner must simply show that ‘infringing activity’ is apparent—pointedly, not ‘*the* infringing activity’ alleged in the complaint.” *Id.* at 12B-58. On Nimmer’s reading of § 512(c), a service provider’s CEO indicating awareness of a news article reporting on its website being a general den of infringing activity may be sufficient to establish a triable issue of fact with respect to red flag knowledge. *See id.* at 12B-59. *See also* 1 Raymond T. Nimmer, Information Law § 4:78.30 (Nov. 2024 update) (“[A]rguably, [red flag knowledge] is a looser standard.”).

Nimmer’s critique centers on the textual differences between § 512(c)(1)(A)(i), the actual knowledge provision, and § 512(c)(1)(A)(ii), the red flag knowledge provision. Specifically, Nimmer focuses on the phrases “the material” and “is infringing” in subsection (i), which do not appear in subsection (ii), to conclude that red flag knowledge “deals with *generalities*” in a way that subsection (i) does not. See 4 Nimmer on Copyright § 12B.04[A][1][b][ii] at 12B-59. Nimmer also cites to the legislative history of the DMCA, which explains that “[t]he important intended objective of [the red flag] standard is to exclude sophisticated ‘pirate’ directories—which refer Internet users to other selected Internet sites where pirate software, books, movies, and music can be downloaded or transmitted—from the safe harbor” to bolster his conclusion that “[i]t would hardly nullify the safe harbor to determine that it is lost when a site caters to a clientele devoted to violating the law.” *Id.* at 12B-57–58 & n.61 (quoting S. Rep. No. 105-190, at 48).

As set out below, we have two questions about Nimmer’s interpretation of the red-flag-knowledge provision but ultimately conclude that we need not resolve them to decide this case. Nevertheless, to provide a full explanation we set out our thoughts about the Nimmer view.

First, we think it noteworthy that the examples Nimmer uses specifically concern an internet service provider offering directory services, i.e., services aggregating an array of links to various other websites for their users to more easily locate and visit

directly. But the service providers in *Viacom* and *UMG*, like YouTube here, offer only video hosting services. Although YouTube users may use the comment functionality on its website to post links to offsite content on the YouTube website that another user could access, YouTube does not gather, review, and then offer a directory of these links to its users as contemplated by the Senate Report and *Nimmer* examples. *See generally In re subpoena of Internet Subscribers of Cox Commc'ns, LLC*, 148 F.4th 1056, 1067 (9th Cir. 2025) (“The plain text of § 512 indicates that the safe harbor for which a service provider qualifies depends on the function the service provider performed with respect to the infringement at issue.”).

Second, the language of § 512(c)(1)(A)(iii) incorporates both the actual knowledge and red flag knowledge standards and assumes either is sufficient to require removal of “the material” so that a service provider may still retain safe-harbor protection. *Nimmer*’s view may be correct in the context of web directories, where a website exclusively dedicated to content piracy might openly advertise itself as such so conspicuously as to necessitate expeditious removal from the directory without requiring knowledge of a specific example of confirmed infringement on the website. But that view is not necessarily accurate in the context of a single website dedicated to video hosting services.

For a video hosting service like YouTube, if general knowledge that the service is being used to occasionally achieve infringement were sufficient to require expeditious removal of all

infringing activity, the service provider would then be left to scour its website for all instances of potential infringement in the hopes of retaining its safe-harbor protection. Even after such scouring, the service provider would still be left not knowing whether it had removed enough material (or the right material) to secure safe-harbor protection through subsection (iii). This is the sort of self-monitoring process that the DMCA emphatically does not require of service providers. See § 512(m)(1) (stating that safe-harbor protection is not conditioned on “a service provider monitoring its service or affirmatively seeking facts indicating infringing activity”); *In re Aimster Copyright Litig.*, 334 F.3d 643, 655 (7th Cir. 2003) (“The common element of [the] safe harbors is that the service provider must do what it can reasonably be asked to do to prevent the use of its service by ‘repeat infringers.’”). Cf. *EMI Christian Music Group, Inc. v. MP3tunes, LLC*, 844 F.3d 79, 91 (2d Cir. 2016) (“Based on the available evidence, a reasonable jury could have concluded that it was reasonable for MP3tunes to track users who repeatedly created links to infringing content in the sideload.com index or who copied files from those links. . . . [D]oing so would not require MP3tunes to ‘monitor’ or ‘affirmatively seek facts’ about infringing activity in a manner inconsistent with § 512(m)(1) because it already had adequate information at its disposal in the form of takedown notices provided by EMI as to which links were allegedly infringing.”) (citations omitted).

Importantly for this case, Nimmer also notes that “[e]ven if [a service provider] receives tens of thousands [of] notifications as to particulars within [a website hosted on the service], all it must

do is to disable access to those precise items named in the umpteen notifications.” 4 Nimmer on Copyright § 12B.04[A][1][d] at 12B-82. So the simple accumulation of adequately-addressed DMCA takedown requests alone cannot amount to red flag knowledge of other infringing activity which precludes safe-harbor protection. *See id.* And “the copyright owner cannot short-circuit the notification procedure through a blanket notification;” the DMCA considers such an attempt a failed notification “and therefore counts it as of no effect.” *Id.* at 12B-82–83. Because Athos bases its red flag knowledge argument on the volume of DMCA takedown requests it has submitted to YouTube, and on YouTube’s refusal to remove additional content not noticed in those takedown requests, *see* Appellant’s Br. at 19–20, Nimmer’s criticism of *Viacom* and *UMG* has no effect on our resolution of the red flag knowledge issue in this case.

2

Athos suggests that YouTube’s video-hash-matching technology was “built to provide actual knowledge of infringement[.]” Appellant’s Br. at 18. But Athos has not advanced any evidence that YouTube has actual knowledge of any infringing material that it does not expeditiously remove. Indeed, Athos appears to concede that YouTube only gains *actual* knowledge of specific infringement once it receives a takedown request identifying the exact location of particular infringing material. *See* Appellant’s Br. at 18 (“YouTube does not remove the clips once it becomes aware of the red flag from which infringing activity is apparent; but rather, waits until it receives ‘actual knowledge’ in the form of another

takedown request.”). We conclude that Athos has not offered any evidence that YouTube ever had actual knowledge of infringement that was not being removed from the site because the parties agree that YouTube expeditiously removes specific videos identified by URL in a DMCA takedown request. *See id.* at 20 (“YouTube chose to remove only the URL of the clip of which it had actual knowledge.”).

The record reflects that the copyright management tools could not have provided YouTube with actual knowledge because they were never used in the manner Athos suggests. As discussed earlier, Athos has not presented evidence that any of the video-hash-matching tools it argues were designed to provide actual knowledge are operated automatically or universally by YouTube. Even if it had made such a showing, the video-hash matches generated by YouTube’s technologies, as we understand them from this record, cannot give YouTube actual knowledge of infringing material because the tools do not perform any analysis to determine whether the hash matches they generate are legally infringing or not.³

Without any kind of additional analysis being performed by the copyright management tools or YouTube employees, YouTube cannot have “actually or ‘subjectively’ [known] of specific infringement,” because the tools at most produce lists of

³ Given the record in this case, we express no views on the application of § 512(c) in a case where a service provider has different functionalities and technologies.

possible infringement. *See Viacom*, 676 F.3d at 31. *See also BWP Media USA, Inc. v. Clarity Digit. Grp., LLC*, 820 F.3d 1175, 1181 (10th Cir. 2016) (explaining that “if the infringing content has merely gone through a screening or automated process, the [internet service provider] will generally benefit from the safe harbor’s protection”). To require YouTube to not only run its copyright management tools in the manner Athos suggests, but to then analyze the fruits of those tools to locate other instances of infringement, would amount to requiring YouTube to “affirmatively [seek] facts indicating infringing activity,” in direct contravention of the balance established by Congress in the DMCA. *See* § 512(m)(1). Athos has not shown that YouTube failed to expeditiously remove any infringing material it had actual knowledge of.

3

Athos maintains that, because of its understanding of YouTube’s technologies, a takedown request regarding specific infringement of a given work inherently necessitates removal of multiple other instances of that work appearing in any non-noticed videos on the YouTube platform. *See* Oral Argument Audio at 12:00–13:47. This is not a plausible or workable reading of the DMCA. Nor, as discussed earlier, is this theory rooted in a correct understanding of how the technologies and functionalities at issue actually work. Again, YouTube does not operate any of its copyright management tools universally or automatically in the manner Athos asserts.

In response to the many takedown requests from Athos, YouTube “fully meets its obligations under the safe harbor by removing the precise material noticed.” 4 Nimmer on Copyright § 12B.04 [A][1][d] at 12B-82. Because Athos offers no evidence other than the accumulation of its takedown requests, which it concedes YouTube has adequately responded to by removing the specific infringing videos identified, *see* Appellant’s Br. at 19–20, its red flag knowledge argument fails. *See Capitol Records, Inc. v. MP3tunes, LLC*, 821 F. Supp. 2d 627, 644 (S.D.N.Y. 2011) (“[T]he DMCA does not place the burden of investigation on the internet service provider. . . . [I]f investigation is required to determine whether material is infringing, then those facts and circumstances are not ‘red flags.’”) (citing § 512(m)(1)).

4

Athos’ final argument with respect to knowledge relates to willful blindness. Athos contends that, because of its technologies and functionalities, YouTube behaves in a way that amounts to intentional avoidance of actual or red flag knowledge of infringing material on the YouTube website. This argument is also unconvincing.

“[W]illful blindness is a form of constructive knowledge for contributory trademark infringement.” *Luxottica Grp., S.p.A. v. Airport Mini Mall, LLC*, 932 F.3d 1303, 1313 (11th Cir. 2019). “Willful blindness occurs when a person suspects wrongdoing and deliberately fail[s] to investigate.” *Id.* (internal quotation marks and citation omitted). We agree with the Second Circuit that, because its

application is limited but not abrogated by § 512(m)(1), the “willful blindness doctrine may be applied, in appropriate circumstances, to demonstrate knowledge or awareness of specific instances of infringement under the DMCA.” *Viacom*, 676 F.3d at 35.

Though a service provider “would not qualify for the safe harbor if it had turned a blind eye to ‘red flags’ of obvious infringement,” S. Rep. No. 105-190, at 48, YouTube did not act this way here. Congress’ focus on the obviously “infringing nature of” material that service providers cannot be willfully blind to reflects the intended “common-sense result” that service providers “not be required to make discriminating judgments about potential copyright infringement.” *Id.* at 48–49.

Like the Second Circuit, “we can see no reason to construe [§ 512(c)] as vitiating the protection of § 512(m) and requiring investigation merely because the service provider learns facts raising a *suspicion* of infringement (as opposed to facts making infringement *obvious*).” *Vimeo I*, 826 F.3d at 98 (emphasis in original). Because the record reflects that YouTube’s copyright management tools are only capable of producing lists of potential infringement, those tools are not capable of producing sufficient knowledge to establish (or to create an issue of fact on) YouTube’s willful blindness to any infringing activity here.

B

We conclude with the alleged right and ability of YouTube to control infringing material it purportedly derives a direct financial benefit from. Agreeing with the Second Circuit’s consideration

of an equivalent issue in *Vimeo II*, we hold that YouTube’s moderation and content management features do not constitute a right and ability to control for purposes of the DMCA. We therefore do not address whether YouTube receives a direct financial benefit from infringing material.

The right and ability to control “requires ‘something more’ than the mere ability to remove or block access to materials on [the service provider’s] website.” *Vimeo II*, 125 F.4th at 423. The service provider must exert “‘substantial influence’ on user activities” or otherwise have “induced the infringing activity.” *Id.* at 423–24.

In *Vimeo II*, the Second Circuit explained that “[c]alling attention to selected videos by giving them a sign of approval or displaying them on a Staff Picks channel (or the contrary, by demoting them) did not restrict the freedom of users to post whatever videos they wished.” *Id.* at 425. It also concluded that, given the “huge number of videos posted by users on Vimeo,” the plaintiffs there had failed to show that Vimeo staff had intervened in more than a “tiny percentage.” *Id.* In its view, “denial of eligibility for the safe harbor based on such noncoercive exercises of control over only a small percentage of postings would undermine, rather than carry out, Congress’[] purposes in establishing the safe harbor.” *Id.* at 426.

We agree with the Second Circuit, and conclude that the same reasoning applies here, where the record reflects only that YouTube can remove material and users from its website, can promote and auto-play select videos algorithmically, and can set

policies for content moderation on the site. The record contains no evidence that YouTube exercises control over any more substantial a portion of the YouTube website than did Vimeo in *Vimeo II*.

None of the identified “noncoercive exercises of control” on the part of YouTube amount to substantial influence over user activity sufficient to establish it had the right and ability to control any infringing material at issue in this case. *See id.* “To interpret this provision as [Athos] argue[s]—to deny [YouTube] access to the safe harbor merely because of the tiny influences it exercised—would subject [YouTube] to a huge expense in monitoring millions of posts to protect itself against the possibility of liability for infringements.” *Id.* at 426–27. Athos’ approach would have us upset the balance between copyright owners and service providers struck by Congress in the DMCA. If that balance is to be recalibrated, it is a matter for Congress.

IV

Aside from specific takedown requests, YouTube’s technologies and functionalities do not produce actual or red flag knowledge of specific infringing material. YouTube, moreover, does not exercise the right and ability to control any of the material on its site for purposes of the DMCA. As a result, YouTube is entitled to § 512(c)’s safe-harbor protection. The district court’s summary judgment order in favor of YouTube is affirmed.

AFFIRMED.