IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF MARYLAND

WIKIMEDIA FOUNDATION, Plaintiff, v. NATIONAL SECURITY AGENCY, *et al.,* Defendants.	Civil Action No. 1:15-cv-00662-TSE

Exhibit 16

IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF MARYLAND

WIKIMEDIA FOUNDATION, Plaintiff, v. NATIONAL SECURITY AGENCY, *et al.*, Defendants.	)	No. 1:15-cv-0662 (TSE)

(THIRD) DECLARATION OF DR. HENNING SCHULZRINNE

Dr. Henning Schulzrinne, for his (third) declaration pursuant to 28 U.S.C. § 1746, deposes and says as follows:

INTRODUCTION AND SUMMARY

1. I am the Julian Clarence Levi Professor of Computer Science at Columbia University in New York, New York. I previously submitted two declarations in this case, dated November 12, 2018 and February 15, 2019. I submit this third declaration at the request of the United States Department of Justice to address the conclusions reached by Mr. Scott Bradner in his reply declaration filed on March 8, 2019, including Mr. Bradner's assessment of conclusions reached in my February 15 declaration. My background and qualifications in the fields of computer science, electrical engineering, and digital communications technology; the sources of information I considered in arriving at the conclusions stated in this case (apart from those cited herein); and my compensation for my services in this matter, are all stated in my prior declarations.

2. For the reasons I detail herein, it remains my conclusion that the hypothesis advanced in this case by plaintiff Wikimedia Foundation ("Wikimedia"), that the National Security Agency ("NSA"), in the course of conducting Upstream collection, must as a matter of

technological necessity be intercepting, copying, and reviewing at least some of Wikimedia's electronic communications that traverse the Internet, is incorrect. Based on what is publicly known about the NSA's Upstream collection technique, the NSA in theory could be conducting this activity, at least as Wikimedia conceives of it, in a number of technically feasible, readily implemented ways that could avoid NSA interaction with Wikimedia's online communications. Nothing stated in Mr. Bradner's reply declaration, or in Wikimedia's sur-reply, alters that conclusion.

3. I also adhere to the conclusions reached in my second declaration, (Second) Declaration of Dr. Henning Schulzrinne (hereinafter, "Second Decl."). In his first declaration Mr. Bradner concluded (i) that the NSA "most likely" uses his copy-all-then-scan configuration to conduct Upstream collection, and (ii) that it is "implausible" that the NSA employs a filter-then-copy-and-scan approach, using whitelisting and blacklisting techniques, such as I described in my first declaration; and (iii) that that even if the NSA uses one or more of the techniques I describe, it is still "virtually certain" that the NSA copies and scans at least some of Wikimedia's communications. *See generally* Declaration of Scott Bradner (hereinafter, "Bradner Decl."). In my second declaration, I explained that none of these conclusions has a foundation in Internet technology or engineering, and instead are based principally on assumptions Mr. Bradner makes about the NSA's practices and priorities, its resources and capabilities, and its Upstream surveillance targets, matters about which Mr. Bradner has no specialized knowledge or information. These conclusions also remain unaltered by Mr. Bradner's reply declaration, or Wikimedia's sur-reply.

4. In his second declaration, Mr. Bradner takes a somewhat different approach. The central thesis of Mr. Bradner's second declaration is that the use of traffic-mirroring techniques (that is, whitelisting and blacklisting) to filter the communications traversing a monitored link before communications of interest are copied and then scanned for selectors, would "conflict[ ]" with "the government's own descriptions" of the Upstream program. Reply Declaration of Scott Bradner (hereinafter "Bradner Reply Decl.") ¶¶ 6, 12, 30. By the "government's own

descriptions," Mr. Bradner means, first, a single statement contained in an 80-page opinion issued by the Foreign Intelligence Surveillance Court ("FISC") in October 2011, regarding the NSA's acquisition of wholly domestic "about" communications, Bradner Reply Decl., App'x P, and second, a single remark in the nearly 200-page report of the Privacy and Civil Liberties Oversight Board on the NSA's Section 702 surveillance program (hereinafter, "PCLOB Section 702 Report"), regarding the NSA's goal of "comprehensively" acquiring communications to or from its targets, Bradner Reply Decl., App'x F.

5. As I already discussed in my prior declaration, the use of a filter-then-copy-and-scan approach to Upstream collection would be entirely consistent with both of these statements. Second Decl. ¶¶ 56-58. The conflict perceived by Mr. Bradner with the FISC's statement arises not from any technical grounds but his own speculative interpretation of that statement, and his assumption that a remark made by the FISC nearly eight years ago still reflects the technical realities of Upstream collection today. The supposed conflict that Mr. Bradner attempts to seize on with the PCLOB's remark ignores the difference between a statement of aspirations on a printed page and the real-world challenges of designing, constructing, deploying, maintaining, and paying for the collection systems required to implement the kind of Upstream collection process that Mr. Bradner envisions. Moreover, he misconstrues the PCLOB's description of the NSA's prior collection of "about" communications as an aim of the program, whereas the PCLOB clearly described "about" collection as "byproduct" of the NSA's efforts to acquire communications sent to or from its Upstream surveillance targets.

6. Mr. Bradner also opines that a filter-then-copy-and-scan approach "conflicts with other technical and practical necessities of conducting" Upstream surveillance. Bradner Reply Decl. ¶¶ 7, 61-112. As I also discuss below, these so-called "technical and practical necessities" are based on the unsupported conclusions Mr. Bradner draws from the FISC and PCLOB statements, as well as many of the same non-technical, speculative assumptions made in his previous declaration about the NSA's surveillance practices and priorities, resources and

capabilities, and its targets, about which Mr. Bradner has no actual knowledge. As such, they still provide no basis in Internet technology or engineering for the conclusions he reaches.

7. Mr. Bradner remarks that I do not offer evidence that the NSA is actually using whitelisting and/or blacklisting techniques in the course of Upstream surveillance that would avoid interaction with Wikimedia's communications. Bradner Reply Decl. ¶¶ 6, 57-58. The same can be said of Mr. Bradner, of course, that he has offered no evidence—only speculation—that the NSA conducts Upstream surveillance in the manner that he describes. As I have explained before, I did not attempt in my previous declarations to reach conclusions about how the NSA actually conducts Upstream surveillance, because I was not asked by the Department of Justice to opine on that question, because the question implicates operational details of Upstream surveillance that remain classified, and because it would have required that I engage in the same sort of speculation as Mr. Bradner, concerning the NSA's actual surveillance practices, capabilities, and targets, about which neither I nor Mr. Bradner has any specialized knowledge or information. *See* Second Decl. ¶ 3. I do not engage in such speculation now, either.

8. As was the case with my first two declarations, in reaching the conclusions stated herein I have not considered, nor have I been provided with, any classified or other non-public information concerning Upstream surveillance.

MR. BRADNER'S CONCLUSION THAT WHITELISTING AND BLACKLISTING WOULD "CONFLICT" WITH THE FISC'S 2011 STATEMENT REGARDING THE COLLECTION OF WHOLLY DOMESTIC "ABOUT" COMMUNICATIONS LACKS A FOUNDATION IN INTERNET TECHNOLOGY AND ENGINEERING

9. As Mr. Bradner observes, in an October 2011 opinion concerning the legal implications of the NSA's collection of so-called multi-communication transactions ("MCTs"), the FISC stated that:

> Indeed, the government readily concedes that NSA will acquire a wholly domestic "about" communication if the transaction containing the communication is routed through an international Internet link being monitored by NSA or is routed through a foreign server. See June 1 Submission at 29.

Bradner Reply Decl., App'x P at 45.

10. In his previous declaration, Mr. Bradner cited this statement in support of a very different conclusion than the one he offers now. Earlier, he cited the FISC's statement as evidence only that the NSA, at least at some monitored links, was not using "an IP filter to eliminate wholly domestic [communications] before" copying and scanning. Bradner Decl. ¶ 294; *see* Second Decl. ¶ 56. Mr. Bradner did not cite this conclusion as a ground for viewing his copy-all-then-scan configuration as "most likely," or a filter-then-copy-and-scan approach as "implausible." *See* Bradner Decl. ¶¶ 282-89, 366-67. Nevertheless Wikimedia, in its summary judgment opposition brief, argued that that the acquisition of some wholly domestic communications, even at a so-called "international Internet link," is inconsistent with the use of the traffic-mirroring techniques I have described, even though Mr. Bradner had not made any such assertion in his first declaration filed at that time. *See* Second Decl. ¶ 56. I explained, therefore, why use of whitelisting and blacklisting techniques would be consistent with the acquisition of at least some wholly domestic "about" communications, Second Decl. ¶¶ 56-58, a conclusion that Mr. Bradner does not dispute, *see* Bradner Reply Decl. ¶¶ 36, 111.

11. Only now, in his second declaration, does Mr. Bradner make an argument that implementing a filter-then-copy-and-scan approach would be inconsistent with the collection of wholly domestic about communications, as referred to in the FISC's opinion. Bradner Reply Decl. ¶¶ 32-45. Indeed, he is emphatic that the FISC's statement would have to be "false" for the NSA to employ a filter-then-copy-and-scan approach. *Id.* ¶ 154(1). This is the case, according to Mr. Bradner, because the use of whitelists or blacklists to filter out certain communications traversing a monitored link could filter out at least some wholly domestic about communications, containing targeted selectors, that the NSA might otherwise acquire. *Id*. ¶ 44. That outcome, Mr. Bradner states, "would not be consistent with the FISC's statement that *all* 'about' communications 'will' be acquired." *Id*. ¶ 42 (emphasis mine).

12. But the FISC did not state that "all" wholly domestic about communications crossing a monitored international Internet link will be acquired by the NSA. It said that "a wholly domestic 'about" communication" will be acquired if it crosses such a link, or is routed through a

foreign server. Mr. Bradner arrives at the conclusion that whitelisting and blacklisting would be inconsistent with the FISC's statement by imputing to the FISC a term, "all," that the FISC did not use. His contention that a filter-then-copy-and-scan approach would not be consistent with the acquisition of wholly domestic about communications is not based on a disagreement over Internet technology or engineering, but on his current reading of the FISC's 2011 statement, one that he did not propose in his first declaration, *see* Bradner Decl. ¶ 296.

13. There are numerous reasons why the statement in the FISC's 2011 opinion does not support Mr. Bradner's conclusion about how Upstream collection must operate today. First, as noted, the FISC simply does not use the term "all" in its statement regarding the collection of wholly domestic about communications.

14. Second, Mr. Bradner does not take into account the context in which the FISC made this remark. The paragraph in which the statement appears concerned the NSA's inability to prevent the acquisition of certain wholly domestic communications, and the Government's suggestion that these acquisitions resulted from a "failure" of the NSA's "technical means." Bradner Reply Decl., App'x P at [45]. The FISC was unwilling to accept that explanation, finding no reason to conclude that the collection of wholly domestic communications was attributable to malfunctions or failures in the NSA's collection equipment. This was the point at which the FISC remarked that the Government had conceded that a wholly domestic about communication would be collected if it crossed a monitored international Internet link, or was routed through a foreign server. In other words, it appears that when the FISC made this remark it was explaining that acquisitions of wholly domestic about communications would occur as the result of technical limitations in the equipment's normal operation (rather than as the result of a malfunction), and was not making a statement about the scope or completeness of those acquisitions.

15. Third, and telling, in an earlier portion of the FISC's opinion where it was discussing the same phenomenon—the NSA's inability to prevent the acquisition of at least some wholly domestic communications—the FISC stated that due to this limitation on the NSA's abilities, the "NSA *may* acquire wholly domestic communications." Bradner Decl., App'x P at 35 n.34

(emphasis mine). Mr. Bradner's conclusion, that according to the FISC the NSA obtained all wholly domestic about communications crossing a monitored link, is thus further undermined by the FISC's clear statement describing the NSA's acquisition of wholly domestic communications as a possibility ("may acquire"), rather than a certainty in all cases.[1]

16. Mr. Bradner also suggests that the FISC must have been "as precise as it possibly could be" when describing the acquisition of wholly domestic about communications, based on the "multiple hearings" it held and the "multiple submissions" it received from the Government before issuing its October 2011 opinion. Bradner Reply Decl. ¶ 38. Of course, this assumption is not a technical basis on which to reach conclusions about how Upstream surveillance is conducted. And I do not understand how Mr. Bradner could know exactly what was said by the Government in those submissions, what was addressed during those hearings, or in what depth, all of which I am advised by the Department of Justice remain classified, in whole or in substantial part. In the end, however, if the FISC meant to say that the NSA would acquire *all* wholly domestic communications crossing a (hypothetically) monitored international Internet link, then the most precise way of expressing that thought would have been to use the word "all." The FISC did not do so.

17. On this subject I observe finally that the FISC's statement cited by Mr. Bradner was made in October 2011, based on a June 1, 2011 submission by the Government, Bradner Reply Decl., App'x P, at 45, nearly eight years ago. Even if the FISC meant that in June 2011 the NSA would acquire all wholly domestic "about" communications crossing an international Internet

[1] Although in this earlier passage the FISC referred to the collection of "wholly domestic communications," rather than "wholly domestic about communications," as in the later passage, there is no reason to believe that it meant to suggest that the collection of wholly domestic communications in general (including both "to/from" and "about" communications) was only a possibility, but at the same time that collection of all wholly domestic "about" communications was a certainty. At least as described in the PCLOB Section 702 Report, the underlying causes for acquisition of wholly domestic communications are similar for both cases: once a communication (for example, an email), makes it past any IP filters, the scanning mechanism would be just as likely to pick up wholly domestic "to" or "from" communications that contain targeted selectors as "about" communications containing targeted selectors. *See* PCLOB Section 702 Report at 10, 37, 84-86, 123, 124, 144-45.

link theoretically monitored by the NSA, and it could be inferred, therefore, that the NSA was not then employing whitelist or blacklist filters, it cannot be taken for granted that the situation has remained the same since then. The growth of Internet traffic, even since 2011, has been enormous. Between 2012 and 2016 alone, the volume of traffic (as measured in used international bandwidth by Telegeography, Inc.) more than quadrupled from 100 terabits per second to over 400 terabits per second[2]. *See* *https://blog.telegeography.com/shaping-the-global-wholesale-bandwidth-market* (figure 1). Given the growth in international Internet traffic since 2011, the possibility cannot be ignored that the NSA at some point might have adopted a form of whitelisting or blacklisting at monitored links to reduce the technical and logistical challenges and costs of processing large volumes of traffic, as discussed in my second declaration, Second Decl. ¶¶ 20-21. Neither Mr. Bradner nor I can know which is the case, of course, but it cannot be taken for granted that the assumption made by Mr. Bradner is the correct one.

18. As I have noted, Mr. Bradner does not disagree with my conclusion that the use of whitelisting and blacklisting would be consistent with the acquisition of at least some wholly domestic "about" communications at international Internet links that hypothetically might be monitored by the NSA. *See* Bradner Reply Decl. ¶¶ 36, 111. Therefore, because it is speculative to suggest (i) that the FISC meant to say that the NSA would collect *all* wholly domestic "about" communications if they crossed a monitored international Internet link in 2011, and (ii) even if that was the FISC's meaning in 2011, that the NSA would continue to collect all wholly domestic about communications at such links years afterward, Mr. Bradner lacks a technical basis on which to conclude that whitelisting and blacklisting would be inconsistent with the FISC's statement.

19. In sum, I find no reason, and certainly none based in Internet technology or engineering, to conclude that whitelisting or blacklisting at international Internet links that the NSA may in theory be monitoring would be inconsistent with the 2011 FISC statement cited by Mr. Bradner.

[2] This traffic volume includes all international traffic, not just traffic entering and exiting the United States. The growth rates, however, have been similar across all major geographies.

THE GOAL OF "COMPREHENSIVELY" ACQUIRING TARGETS' COMMUNICATIONS IS NOT A TECHNICAL BASIS ON WHICH TO CONCLUDE THAT THE NSA DOES NOT EMPLOY WHITELISTING OR BLACKLISTING TECHNIQUES

20. Mr. Bradner next opines that a filter-then-copy-and-scan approach to Upstream collection process would be "incompatible" with the goal of "comprehensively acquir[ing] communications that are sent to or from [the NSA's] targets," as stated on pages 10 and 123 of the PCLOB Section 702 Report. Bradner Reply Decl. ¶¶ 46-54.[3] He goes so far as to say that the PCLOB's statement would have to be "false" for the NSA to conduct Upstream collection using the traffic-mirroring techniques I have described. *ID.* ¶ 154(2). As I have explained previously, Mr. Bradner offers no technical basis for concluding as a matter of concrete reality that the NSA "must be" copying and reviewing all communications crossing a monitored link, Bradner Reply Decl. ¶ 29, simply because the NSA might wish in the abstract to obtain as many of its targets' communications as it can. Second Decl. ¶¶ 71-75.

21. Specifically, I explained that one cannot assume away, based on a stated aim of comprehensiveness, the many technical, logistical, and financial hurdles, and competing mission priorities, that would stand in the way of designing, constructing, deploying, and maintaining the kind of collection systems envisioned by Mr. Bradner. Second Decl. ¶ 73. (See also paragraphs 20-21 of my second declaration, in which I describe some of the technical and logistical challenges of implementing Mr. Bradner's copy-all-then-scan approach.) Therefore, one cannot draw meaningful conclusions about the technical details of the NSA's Upstream collection systems with nothing more to go on than an abstract goal of comprehensiveness. Second Decl. ¶ 74.

22. Mr. Bradner does not dispute these practical realities. He himself states that "the NSA must operate in the real world and deal with the technical and operational limitations

[3] I note that in contrast the PCLOB describes the collection of "about" communications as a "byproduct" of the NSA's efforts to collect communications sent to or from its targets. PCLOB Section 702 Report at 10, 123. The PCLOB's description of "about" collection as a "byproduct" rather than an objective of Upstream collection tends to rebut Mr. Bradner's interpretation of the FISC's October 2011 statement to mean that the NSA must have configured Upstream collection in such a way as to acquire "all" wholly domestic about communications on a monitored link. *See also* paragraphs 47-49, below.

inherent in the Internet." Bradner Reply Decl. ¶ 10. Yet he continues to disregard these same realities when he attempts to draw technical conclusions about how Upstream surveillance "must be" conducted from a single statement by the PCLOB attributing a goal of "comprehensiveness" to the NSA.

23. Mr. Bradner remarks that the use of whitelists and blacklists to reduce the technical and logistical difficulties and costs of Upstream collection is "incompatible" with the goal of completeness, Bradner Reply Decl. ¶ 51, but they are entirely compatible with reducing the formidable technical and logistical burdens and costs of processing large volumes of communications traffic, and Mr. Bradner has no way of knowing, and certainly no technical way of determining, whether such "real world" constraints have convinced or compelled the NSA to compromise the goal of completeness. It is, at bottom, the hard realities of achieving Mr. Bradner's vision of Upstream collection that may have proven incompatible with that goal, just as they are incompatible with a bare assumption that the NSA has succeeded in implementing that vision simply because it wants to.

24. Mr. Bradner opines that the single term "comprehensively" provides an "appropriate basis on which to explain the technological implementation" of Upstream collection, because "the PCLOB used the term 'comprehensively' to explain the need for the NSA's specific technological implementation" of the program. Bradner Reply Decl. ¶ 48. But the PCLOB's Section 702 Report, being an unclassified public report about a highly classified foreign-intelligence gathering activity, contains no specific technical detail about how Upstream collection is conducted. It certainly includes no technical detail at the level required to support the conclusions reached about Upstream collection by Mr. Bradner, except as a matter of speculation and conjecture.

25. "If wishes were horses," as they say, then one could simply assume that the NSA had the technical, logistical, and financial wherewithal, consistent with all its other mission requirements, to deploy an Upstream collection architecture capable of comprehensively acquiring every single one of its targets' online communications, just because it would like to do

so. And perhaps it has succeeded in doing so. I (like Mr. Bradner) do not know. But the assumption has no basis in the "real world" of Internet technology and engineering, Bradner Reply Decl. ¶ 10, in which the NSA must operate.

MR. BRADNER'S OPINION THAT WHITELISTING AND/OR BLACKLISTING WOULD "CONFLICT" WITH "TECHNICAL AND PRACTICAL NECESSITIES" OF UPSTREAM COLLECTION RESTS ON SPECULATION RATHER THAN A BASIS IN INTERNET TECHNOLOGY AND ENGINEERING

26. In addition to the inferences he draws from his interpretation of the FISC's 2011 statement, and the PCLOB's use of the term "comprehensive" to describe Upstream's objective, Mr. Bradner also attempts to reach conclusions based on what he calls "technical and practical necessities" that make clear, in his view, that the NSA is copy and scanning at least some of Wikimedia's communications. As I discuss below, these "technical and practical necessities" turn out again to be assumptions by Mr. Bradner about the NSA's surveillance practices and priorities, resources and capabilities, and the nature and behavior of its Upstream surveillance targets, mixed with the unsupported inferences he draws from the FISC and PCLOB statements discussed above. They supply no basis in Internet technology and engineering for concluding that the NSA "most likely" uses a copy-all-then-scan configuration to conduct Upstream surveillance, or that a filter-then-copy-and-scan approach is "implausible."

Whitelisting IP Addresses of Interest

27. Developing and maintaining whitelists: In Mr. Bradner's view whitelisting would be unworkable for purposes of Upstream collection because it would require "knowing in advance all of the IP addresses that might be used by each of the NSA's targets," Bradner Reply Decl. ¶ 68, and maintaining "comprehensive" information on where they will be, what sites they communicate with, and the protocols they use, Bradner Reply Decl. ¶ 69. *See also id.* ¶ 88. These conclusions lack a non-speculative technical basis.

28. As a technical matter, the NSA would not need to know all of its targets' IP addresses or gather comprehensive information on their whereabouts, the sites they visit, or the types of online communications in which they engage, before it could whitelist the IP addresses that the agency already knows about. As pointed out in the Government's reply brief, the NSA

could acquire information about its targets' IP addresses from communications acquired under the Section 702 PRISM program (*see* PCLOB Section 702 Report (Bradner Decl., App'x F) at 7, 33-34), Executive Order 12,333, prior Upstream acquisitions, information obtained from other U.S. intelligence agencies, and other sources. Reply Brief in Support of Defendants' Motion for Summary Judgment at 12-13. Mr. Bradner does not dispute this. Presumably communications acquired from these sources would also provide insight into particular websites of interest that targets visit, and their preferred modes of communication, thus allowing for further enhancement of the NSA's whitelist(s).

29. Mr. Bradner suggests no technological reason why it would not be possible for the NSA to conduct Upstream surveillance in this fashion. Doing so would not mean that the NSA "is only interested in the people and processes it already knows about and that it has decided to actively ignore everything else." Bradner Reply Decl. ¶ 89. Nothing about the use of whitelists would prevent the NSA from learning about new potential targets of interest, or new information about the communications of its existing targets, either from ongoing Upstream collection or other intelligence sources (in Mr. Bradner's words, discovering which streetlights to look under for your keys, *id.* ¶ 70), and then updating its whitelists accordingly.

30. There would be no need whatsoever for the NSA to know in advance all the information Mr. Bradner refers to unless one works backward from the conclusion, as Mr. Bradner does, that the NSA's acquisition of its targets' communications must be comprehensive, *see* Bradner Reply Decl. ¶¶ 69, 88, based on the PCLOB's passing remark about the NSA's goals. As I have discussed, the PCLOB's report does not provide a technical justification for concluding that the NSA's acquisition of its targets' communications is comprehensive in fact, or drawing inferences based on that premise about how Upstream surveillance is conducted.

31. Mr. Bradner observes that whitelists would have to be updated as targets were added or removed, or changed their locations or modes of communications. Bradner Reply Decl. ¶ 85. This is true, of course, but how often that would have to occur Mr. Bradner does not say,

because he cannot know, and he certainly offers no reason why updating whitelists as needed would be beyond the NSA's capabilities.

32. Target numerosity: Mr. Bradner next repeats his argument that the NSA's targets are too numerous to make the development and maintenance of whitelists of their IP addresses practical. Bradner Reply Decl. ¶¶ 75-76; *see* Bradner Decl. ¶ 366(d). I have already explained that in so arguing Mr. Bradner is making speculative assumptions about the number, nature, and communications habits of the NSA's Upstream targets, about which Mr. Bradner has no information. First Decl. ¶¶ 45-48. He does not maintain otherwise. Instead, he offers a hypothetical in place of facts: if the NSA had 1,000 Upstream targets in 2011, when according to the FISC it acquired 26 million communications using its Upstream collection technique, then on average each target would have had to engage in at least 26,000 online communications that year. Bradner Reply Decl. ¶ 76. But if we used another randomly chosen number of Upstream targets, say 5,000, then the number of communications each target would have had to send or receive in 2011 drops to 5,200, or approximately 14 per day. That is a trivial number considering (i) that "communications" can include not only such media as email but also all the individual HTTPS (or HTTP) requests and responses made during a single visit to a website, and (ii) that the NSA's Upstream targets could include organizations as well as individuals, organizations employing dozens or hundreds of persons capable of generating thousands of communications per day. Mr. Bradner has given no factual, technical basis for concluding that the number of the NSA's Upstream targets would make whitelisting unworkable.

33. Target mobility: Mr. Bradner also repeats his arguments that NSA use of whitelists is not "remotely possible" because its targets may move around, and as a result their IP addresses could change; because targets may use intermediary communications services, such as virtual private networks (VPNs, which remove the targets' IP address from packet headers during certain legs of their journeys across the Internet); and because targets may use multiple Internet service providers (ISPs), that assign different IP addresses to their subscribers. Bradner Reply Decl. ¶¶ 77-78, 82. In my view, Mr. Bradner overstates the matter. To take just one example, if the

NSA were seeking to track a targeted individual's email, determining the source and destination IP addresses to whitelist could be automated readily. An exhaustive list of SMTP destination IP addresses for any email address can be looked up online via DNS (MX records), and the list of source IP addresses can be derived from the DNS SPF entries. These IP addresses would not be affected by user mobility, and would only depend on the sender or receiver email address, not the user's current location.

34. But again, most fundamentally, Mr. Bradner is making assumptions about the nature, mobility, and communications practices of the NSA's targets. Only the NSA knows the extent to which its targets' IP addresses change due to their mobility, the extent to which they use VPNs, or multiple ISPs, to communicate, and therefore whether whitelisting would be impractical for purposes of meeting its intelligence-collection needs. Whitelisting could not be considered technologically impossible, however, unless one started from the premise that the NSA must be "comprehensively" acquiring every single one of its targets' communications without fail, and worked backward from there. As I have explained, Mr. Bradner has offered no justification, certainly no technological justification, for adopting that premise.

35. Further, as I note in my second declaration, to the extent a target moves from place to place within a given geographic area, the NSA could whitelist a set of IP addresses, rather than just a single address, associated with the geographic region where the target is believed to be located. Second Decl. ¶ 47. Mr. Bradner does not dispute this, but points out that using a range of IP addresses could increase the chances that communications to or from Wikimedia could be copied and scanned. Bradner Reply Decl. ¶¶ 80, 82. This is true in the abstract, but Mr. Bradner cites no information about the extent to which the NSA might find it necessary to whitelist ranges of IP addresses—how many, how large, in what geographic areas—in order to reliably monitor its targets' communications. As a result he can only speculate whether whitelisting ranges of IP addresses would result in the NSA copying and scanning Wikimedia communications.

36. Combined whitelisting and blacklisting to selectively acquire web communications: I previously observed that if it wished the NSA could, as a technological matter, simultaneously whitelist the IP addresses of particular websites, webmail services, and/or chatrooms of interest while blacklisting all other HTTP and HTTPS traffic, and thus obtain access to web communications of interest without necessarily copying and scanning Wikimedia's. Second Decl. ¶¶ 35, 36(b), 37. In response, Mr. Bradner remarks that "many" websites and other web-based services are now making use of content distribution networks ("CDNs"), which can have different and changing IP addresses around the world. Bradner Reply Decl. ¶ 84. Yet again, the extent to which websites, webmail services, and chatrooms of interest to the NSA (if any) are using CDNs for their communications, and the extent to which their use of CDNs would make it difficult or unworkable for the NSA to use whitelists to track its targets' communications, are matters known to the NSA, but about which Mr. Bradner has no information, and can only speculate.

37. Whitelisting by protocol: Mr. Bradner also points out (i) that if the NSA used protocol-based whitelisting (as opposed to whitelisting by IP addresses), it could miss communications that it ideally it might want to review, Bradner Reply Decl. ¶ 72, and (ii) that the NSA could not in fact be using whitelists to exclude all communications using web protocols (HTTP and HTTPS) while at the same time collecting at least some web communications, *id.* ¶ 73. These are points that Mr. Bradner has raised before, Bradner Decl. ¶ 366(f), (g), and that I addressed in my second declaration, Second Decl. ¶¶ 34-35, 36(b), 37. For the reasons already explained in my second declaration, these observations by Mr. Bradner supply no basis in Internet technology or engineering for concluding that whitelisting by the NSA would be implausible.

38. I do not mean to suggest that developing and maintaining whitelists for purposes of NSA Upstream collection would necessarily be "easy," or that "the NSA could get by with a very simple set of whitelist rules." *See* Bradner Reply Decl. ¶¶ 70, 85, 88. But neither I nor Mr. Bradner can know how difficult or easy it would be without far more detailed information about the number, nature, and communications habits of the NSA's Upstream targets. I do not

presume, as does Mr. Bradner, that the NSA would be incapable of developing and maintaining sufficiently reliable whitelists to meet its intelligence needs. Mr. Bradner has presented no basis in Internet technology or engineering for so assuming.

Blacklisting by IP Address or Protocol

39. Mr. Bradner states that "[t]here are multiple reasons" why the use of blacklists to filter out communications before they are copied or scanned in the Upstream process "would be incompatible with the public descriptions of the NSA's upstream collection program," Bradner Reply Decl. ¶ 92, but then proceeds to offer none. Instead he contends (i) that blacklisting Wikimedia IP addresses so as to prevent copying and scanning Wikimedia communications is "improbable," and (ii) that blacklisting Wikimedia IP addresses would not "guarantee" that Wikimedia communications are not copied and scanned during the Upstream collection process. *Id.* ¶¶ 93-101. Both are points that Mr. Bradner has raised before, that I have already addressed, and that still do not constitute non-speculative, technical grounds for deeming it improbable or implausible that the NSA might blacklist Wikimedia communications.

40. Blacklisting high-volume websites, including Wikimedia's: Following the observations in my first declaration that the NSA could blacklist Wikimedia IP addresses to prevent communications to and from Wikimedia from being copied and scanned, First Decl. ¶¶ 78-87, Mr. Bradner responded that he found it "basically inconceivable" and "totally unbelievable" that the NSA would sift through millions of websites to decide which to monitor and which not. Bradner Decl. ¶ 367(a). I then explained, in my second declaration, that I had no such extreme measure in mind, but rather, the trivial task of creating a blacklist of numerous high-volume popular websites of potentially low interest to the NSA, such as Wikimedia's, to eliminate unwanted volumes of communications that would otherwise have to be processed. Second Decl. ¶¶ 39-41.

41. Now Mr. Bradner takes the opposite tack, arguing that it is "very unlikely" that the NSA would decide to "specifically blacklist Wikimedia communications to reduce the load" on its collection apparatus, in light of the relatively small percentage of inter-regional Internet capacity

that Wikimedia communications traffic represents. Bradner Reply Decl. ¶ 96. I note first that a more enlightening statistic would be the percentage of the NSA's processing capacity that Wikimedia traffic represents, but Mr. Bradner has no way of knowing that information any more than he could know, rather than speculate, whether the NSA might find reason to blacklist Wikimedia communications or not. But it is also important to note that Mr. Bradner has once again missed the point of my observations. I did not suggest, in my second declaration, a scenario in which the NSA specifically singles out Wikimedia sites for blacklisting. I pointed out that Wikimedia's websites would naturally fall on any blacklist of high-volume, popular websites, of perhaps low interest to the NSA, that the NSA might assemble in order reduce (potentially by as much as 90 percent or more) the technological, logistical, and financial burdens of processing large volumes of unwanted web traffic. Second Decl. ¶ 41. Whether or not the NSA actually does so I do not know, but Mr. Bradner offers no technological basis for dismissing the possibility as improbable.

42. Hypothetical copying and scanning of blacklisted Wikimedia communications: Second, Mr. Bradner again draws attention to three hypothetical scenarios in which communications to or from Wikimedia would be copied and scanned during the Upstream collection process, even if the NSA had blacklisted communications containing Wikimedia IP addresses. Bradner Reply Decl. ¶¶ 97-101; *see* Bradner Decl. ¶ 367(a). These involved transmission of a Wikimedia communication within a multi-communication transaction (MCT) across an international Internet link, transmission of an email to Wikimedia from abroad using a U.S.-based email service, and visits to Wikimedia websites from abroad using a U.S.-based VPN service. I explained in my prior declaration that in each of these scenarios four (or in the third scenario, five) conditions would have to be met before a communication to or from Wikimedia would be copied or scanned, in each case rendering that possibility a matter of speculation. Second Decl. ¶¶ 77-85.

43. Regarding the initial three (or four) conditions required before each scenario could come to pass, Mr. Bradner summarily asserts that these conditions "would likely be frequently

met," yet offers no supporting data or explanation to demonstrate why that would be so. Instead, he focuses his attention principally on the last condition that would have to be met in each scenario before a Wikimedia communication could be copied or scanned—that the communication not be blacklisted for other reasons. Mr. Bradner dismisses as "far-fetched" the possibility that communications of the kind he posited in his three scenarios would be blacklisted, based on conjecture about the value that the NSA might attach to them. Bradner Reply Decl. ¶¶ 99-101. But he does not address the reasons given in my prior declaration to expect that the communications he describes likely would be encrypted, and perhaps blacklisted by the NSA, therefore, if it lacked the ability to decipher them. Second Decl. ¶¶ 80, 82, 84. In any event, if the initial three (or four) conditions in each scenario are not met, whether or not the final condition is met would be of no consequence. And as noted above, Mr. Bradner gives no reason to expect that the occurrence of the initial three (or four) conditions in each scenario would be anything but speculative.

Additional Points Concerning Whitelisting and Blacklisting

44. Mr. Bradner completes his discussion of whitelisting, blacklisting, and the filter-then-copy-and-scan approach generally with additional points, most already made in his prior declaration that I will now address.

45. "Blind spots": Mr. Bradner returns to his earlier remark that if the NSA blacklisted particular types of communications by port or protocol number, then doing so would leave "blind spots" in its collection that "[s]ophisticated" targets could "easily probe" to discover and evade collection of their communications. Bradner Reply Decl. ¶¶ 103-07; *see* Bradner Decl. ¶ 366(b), (e). On this point I observed previously that Mr. Bradner had not explained what targets could "probe," or how, to discover these so-called blind spots, the level of sophistication required, or on what basis he presumed that the NSA's Upstream targets possess the needed sophistication. Second Decl. ¶ 32. At bottom, I further observed, Mr. Bradner could only speculate whether the creation of "blind spots" would be of such genuine concern to the NSA as to dissuade it from utilizing whitelisting or blacklisting techniques. *Id.* ¶ 33.

46. In his second declaration, Mr. Bradner adds to his earlier conjecture by supposing that unspecified foreign intelligence services that may or may not be Upstream targets could test a protocol (type of communication) they suspect the NSA is not monitoring by communicating "actionable" intelligence ("such as the identity of a foreign agent") over that protocol and then observing whether any responsive action is taken. Bradner Reply Decl. ¶ 103. He vaguely suggests that whitelists or blacklists shared with telecommunication service providers could be at unspecified risk of "hacking," without considering additional security measures that could be taken to mitigate those risks. *Id*. ¶ 105. Typically, however, carriers use segregated networks to manage their routers and switches and encrypt network management information. There have been no indications that I am aware of that such carrier management networks have been breached. In short, neither of these suggestions constitutes a non-speculative basis in Internet technology or engineering to support Mr. Bradner's conclusions. And he overlooks the fundamental point I made previously, that only the NSA knows whether it considers these to be genuine risks that it would be unprepared to take in order to implement a filter-then-copy-and-scan approach to Upstream collection. Second Decl. ¶ 33.

47. "About" communications: In his second declaration Mr. Bradner suggests for the first time that whitelisting or blacklisting would be "entirely inconsistent" with the NSA's now-discontinued acquisition of "about" communications. Bradner Reply Decl. ¶ 108; *see also id.* ¶ 78. (As I noted in my second declaration, this was an argument advanced by Wikimedia in its legal brief, but not by Mr. Bradner in his first declaration. Second Decl. ¶ 49.) The position now taken by Mr. Bradner is based on a complete misunderstanding of "about" collection as described in the PCLOB Section 702 report.

48. Mr. Bradner acknowledges, as I have explained, that even if the NSA were to employ whitelists and/or blacklists to implement a filter-then-copy-and-scan approach to Upstream collection, it would still be possible to acquire at least some "about" communications, that is, communications neither to nor from a target but which refer to a selector (e.g., an email address) associated with the target. *See* Second Decl. ¶ 51; Bradner Reply Decl. ¶ 111. Mr.

Bradner's point is that the use of whitelists or blacklists would be incompatible with the *comprehensive* collection of "about" communications. For example, he states (twice) that developing and maintaining whitelists for acquisition of "about" communications "would be impossible to do . . . for a program *meant to capture* the communications of unknown non-targets about targets." Bradner Reply Decl. ¶¶ 69-70 (emphasis mine). In the same vein, he later reasons that "to set up [a] whitelist filter the NSA would have to know in advance which non-targets' IP addresses to whitelist … *in order to find the 'about' communications.*" Bradner Reply Decl. ¶ 110 (emphasis mine). He ends with the conclusion that, using a filter-then-copy-and-scan model, "the only way the NSA *could reliably capture* about communications would be to whitelist all non-wholly domestic communications." Bradner Reply Decl. ¶ 112 (emphasis mine).

49. The premise of Mr. Bradner's argument, however, is incorrect. At least as publicly described by the PCLOB Section 702 Report, in the very passage reproduced by Mr. Bradner in his declaration, the NSA's collection of "about" communications was

> an inevitable *byproduct* of the government's efforts to comprehensively acquire communications that are sent to or from its targets.

Bradner Reply Decl. ¶ 48 (quoting PCLOB Section 702 Report at 10) (emphasis mine). That is to say, while the goal of Upstream collection as described by the PCLOB is to comprehensively acquire communications sent "to or from" the NSA's foreign-intelligence targets, the program was not likewise "meant to capture," Bradner Reply Decl. ¶¶ 69-70, "reliably" find, *id.* ¶¶ 110, 112, or otherwise "comprehensively acquire" "about" communications. *See also* PCLOB Report at 84-86, 123, 124, 144, 145. The incidental collection of only some "about" communications, even if the NSA were using a filter-then-copy-and-scan configuration as described in my second declaration, would be entirely consistent with the PCLOB's description of "about" collection as a byproduct rather than an objective of Upstream collection. For example, if a whitelist included a specific email server, the NSA could scan for targeted email addresses in the whitelisted communications, and would, as described in the PCLOB Section 702 Report, occasionally acquire "about" communications that referred to the targeted email address in addition to

communications to or from the email address. Thus, the "about" communication would be captured because of the way the scanning is implemented. And being incidental, "about" collection would not have required advance knowledge of all non-targets' IP addresses.

50. Collection of web communications, including encrypted HTTPS communications: Mr. Bradner again addresses the reasons why he believes the NSA is collecting web (i.e., HTTP and/or HTTPS) communications, Bradner Reply Decl. ¶¶ 130-36, 154(3)(c)(i); *see* Bradner Decl. ¶¶ 314-15m 366(f), focusing in particular on his earlier opinions (i) that blacklisting the HTTP and HTTPS protocols to prevent the copying and scanning of all web communications "would leave a very large hole in the NSA's coverage." Bradner Reply Decl. ¶ 135; *see* Bradner Decl. ¶ 366(f), (g), and (ii) that encrypted communications collected by the NSA (if any) include, specifically, HTTPS communications, Bradner Reply Decl. ¶¶ 137-39; *see* Bradner Decl. ¶¶ 325, 366(g).[4] Mr. Bradner largely ignores the observation, in my second declaration, that the NSA could use a combined whitelisting/blacklisting technique to block NSA access to all HTTP and HTTPS communications except those to or from IP addresses included on a whitelist containing the addresses of websites, chatrooms, and/or webmail services of intelligence interest. In this way the NSA could obtain access to HTTP and HTTPS communications of interest, while excluding all others, including, hypothetically, Wikimedia's. Second Decl. ¶¶ 35, 36(b), 37.

51. Mr. Bradner opines that this technique "would also leave very large holes in the NSA's coverage," Bradner Reply Decl. ¶ 136; *see also id.* ¶ 139, but whether the "holes" would be so unacceptably large as to motivate the NSA to copy and scan all HTTP and HTTPS communications (including, therefore, Wikimedia's) is a matter implicating the NSA's surveillance priorities, resources, and capabilities, about which Mr. Bradner has no information and can only

[4] I note in passing that Mr. Bradner expresses skepticism at my suggestion that the NSA could collect encrypted communications under its Section 702 authority using its PRISM acquisition method, as opposed to Upstream collection, because providers assisting in PRISM collection "will frequently have direct access to the user's unencrypted communications." Bradner Reply Decl. ¶ 139. But numerous online file-storage and webmail services now store their users' files and/or messages in encrypted formats, either by default or upon request.

speculate. Mr. Bradner also suggests that combined whitelisting and blacklisting of HTTP and HTTPS communications "would also be contrary to the aim of the 'about' collection program." Bradner Reply Decl. ¶ 136. As I discussed, however, in paragraph 49, above, there is no evidence that the NSA ever engaged in an "'about' collection program." The PCLOB Section 702 Report does not state that "about" collection was ever an "aim" of Upstream collection, merely the "byproduct" of efforts to collect communications to and from designated targets.

52. Relative complexity of the copy-all-then-scan and filter-then-copy-and-scan approaches: Mr. Bradner describes his copy-all-then-scan approach as a "simpl[er], mo[re] reliable, and easi[er] to operate architecture" than a filter-then-copy-and-scan configuration, Bradner Reply Decl. ¶ 116, because in his view, "the need to constantly reconfigure the [provider's router or switch] with updated blacklists and whitelists would create the risk of misconfiguration or overloading." *Id.* ¶ 118; *see also id*. ¶ 124.

53. To begin, it is not unusual for a provider to reconfigure its routers and switches on a routine basis to meet the evolving data transmission needs of a dynamic commercial customer base. Whether it would require more frequent reconfiguration to update an NSA whitelist or blacklist is a matter of conjecture that would depend on how frequently the NSA requested such updates, based on how often it makes additions to or deletions from its target list, and how often its targets change their modes of communication. It is also a matter of speculation, as I have observed previously, whether the NSA's targets (or, more precisely, their associated IP addresses), are so numerous that loading a whitelist would run the risk of overloading a router's or switch's processing capacity. Second Decl. ¶ 25. These are all matters about which Mr. Bradner apparently has no information. Moreover, carriers usually implement change-management protocols in which new router configurations are tested in laboratory settings before they are loaded, in order to mitigate any risk of misconfiguration or overloading.

54. The question Mr. Bradner also leaves unanswered is whether the risks he cites would outweigh the daunting technical and logistical difficulties and financial burdens that would complicate implementation of his "simpl[er]" and "easi[er]" approach, as discussed in paragraphs

20-21 of my second declaration. Mr. Bradner does not take issue with my description of these hurdles (with the exception of the marginal observation that an opto-electronic device is not an "esoteric" piece of equipment). Bradner Reply Decl. ¶ 119. When considered from a broader perspective, the asserted simplicity of Mr. Bradner's copy-all-then-scan configuration as compared to a filter-then-copy-and-scan approach becomes less apparent, to say the least, and which way the scales tip in the NSA's eyes is a matter about which Mr. Bradner, once again, can only speculate.

55. Sharing sensitive information with a provider: Mr. Bradner also repeats a point made in his prior declaration that, "in [his] opinion," the NSA would not want to implement a filter-then-copy-and-scan approach because it would require sharing with provider personnel the sensitive information about NSA targets and the scope of its surveillance that would be contained in whitelists and blacklists. Bradner Reply Decl. ¶ 126. This is a point I have already addressed. Second Decl. ¶ 18. Mr. Bradner acknowledges, as I observed previously, that the NSA shares information about its targets with Internet service providers in order to conduct PRISM collection, but suggests that in his estimation the information that would have to be shared with a provider in a whitelist or blacklist is even more sensitive. Bradner Reply Decl. ¶ 126. Whether or not that is so (a matter about which Mr. Bradner has no apparent expertise), it remains the case that the extent to which the NSA would be willing (or find it necessary) to share classified information with an assisting provider in order to conduct Upstream surveillance is a matter about which Mr. Bradner has no specialized knowledge or information.

56. Proposed "channel mirroring": Mr. Bradner also makes a passing reference to a configuration not previously proposed by him, in which a provider would configure its router or switch to copy all communications just on particular channels (circuits), designated by the NSA, that cross a monitored link. Bradner Reply Decl. ¶ 121. He describes that configuration, too, as very simple, static, and involving fewer disclosures of sensitive information to non-government personnel, *id.*, but it is not so simple a picture as Mr. Bradner paints.

57. In brief, due to carrier practices including link aggregation, among others, it is unclear how the NSA could know which optical channels would be carrying the communications of potential targets, at least without a deep understanding of the assisting carrier's routing architecture and configuration. In addition, because of a phenomenon referred to as traffic failover, traffic can move from one logical link to another, so this configuration is unlikely to be stable and unvarying. As a result, the NSA would have to convey to the carrier which IP addresses it would like to monitor, so that the carrier can map these addresses onto logical links and their optical channels. Thus, the NSA would have to convey just as much information to the carrier as with a filter-then-copy-and-scan configuration. Furthermore, the mirroring capability of at least some common routers and switches is limited to a small number of interfaces. Thus, only a small fraction of the router input or output ports could be monitored at any time if all the traffic, unfiltered, is to be copied to these ports.

58. U.K. Section 8(4) collection: Mr. Bradner returns to this subject to prove a point that I did not take great issue with in my second declaration: that the brief filed by the U.K. government in the European Court of Human Rights describing its "Section 8(4)" collection program includes references to "intercept[ing]"--which could be taken to mean copying—"the entire contents of a [circuit]" before the communications stream is filtered, and the remainder then scanned for targets' communications. Bradner Reply Decl. ¶ 140-48; *see* Second Decl. ¶¶ 61-63.[5] Notably, Mr. Bradner does not take issue with the conclusions I previously drew from the U.K. governments' filings: (i) that the U.K. government documents describe a process of filtering to winnow out communications deemed to lack significant intelligence value before communications are scanned for targets' selectors, Second Decl. ¶ 62, and (ii) that even if all communications on a circuit were copied first, the copying and initial filtering could be conducted

[5] Oddly, Mr. Bradner appears to find fault that I based my discussion of the Section 8(4) collection program on one of the same documents he relied on his declaration (the U.K. government brief), and that I failed to cite a document, *not* appended to his first declaration, that he himself did not cite previously, and that he himself now cites only for the first time in his reply declaration. *See* Bradner Reply Decl. ¶ 144; Second Decl. ¶ 61 (citing Bradner Decl., App'x EE).

by the service provider, so that only those communications meeting the filter criteria, rather than all communications, would pass into the government's (whether the GCHQ's, or hypothetically, the NSA's) control. Second Decl. ¶ 64.[6]

THE SCENARIOS ENVISIONED BY MR. BRADNER IN WHICH AT LEAST SOME WIKIMEDIA COMMUNICATIONS WOULD BE COPIED AND SCANNED, EVEN IF THE NSA EMPLOYED TRAFFIC-MIRRORING TECHNIQUES, ARE SPECULATIVE AND CONJECTURAL

59. Finally, I address Mr. Bradner's repeated contention that the filter-then-copy-and-scan configuration using whitelists and/or blacklists would not "guarantee" that the NSA avoids all interaction with Wikimedia communications during the Upstream collection process. Bradner Reply Decl. ¶ 57; *see also id.* ¶¶ 30, 64, 97-101, 115, 154(5). In support of this contention, Mr. Bradner identifies a number of hypothetical scenarios in which Wikimedia communications would be copied and scanned if all of the necessary conditions were met for these scenarios to come to pass. I have already discussed most of the scenarios outlined by Mr. Bradner above, and in my second declaration. I address the scenarios newly conceived of by Mr. Bradner below. Whether any of them has occurred or would ever come to pass is a matter of speculation for which Mr. Bradner gives no evidence.

60. Whitelisting by IP address: First, Mr. Bradner points to a situation in which a user of a whitelisted IP address communicates with Wikimedia, and the communication traverses an international Internet link (hypothetically) monitored by the NSA. Bradner Reply Decl. ¶ 154(5)(a)(i); *see also id.* ¶ 154(5)(b)(iii). I acknowledged in both my first and second declarations the theoretical possibility that Wikimedia communications of this kind could be copied and scanned during Upstream surveillance, even if the NSA used a whitelisting technique

[6] Mr. Bradner wonders why I also addressed the U.S. Government's cyber-defense system known as Einstein 2.0, when he mentioned it only "in passing" in his first declaration. Bradner Reply Decl. ¶ 150. As I explained in my second declaration, I addressed the implications of Einstein 2.0 because Wikimedia, in its legal brief, attempted to rely on Einstein 2.0 as "corroboration" for Mr. Bradner's conclusions, even though Mr. Bradner himself had not done so, Second Decl. ¶ 66, and does not do so now. For all intents and purpose, Mr. Bradner now disregards my explanation of the reasons why conclusions about the Upstream collection process cannot be drawn from Einstein 2.0. *Compare* Second Decl. ¶¶ 68-69 to Bradner Reply Decl. ¶¶ 151-53.

that otherwise excluded Wikimedia's communications. Second Decl. ¶ 43; First Decl. ¶ 81. But this scenario, like the blacklisting scenarios posited by Mr. Bradner, is conjectural. *See* paragraphs 42-43, above; Second Decl. ¶¶ 78-85. Mr. Bradner cites no evidence of the number or geographic locations of persons using whitelisted IP addresses who communicate with Wikimedia, and of course could not do so without knowing the composition of whitelists (hypothetically) employed by the NSA. There is no basis, therefore, on which to conclude that communications between Wikimedia and persons using whitelisted IP addresses would cross every international Internet link to and from the United States (as Wikimedia claims of its communications generally), or, for that matter, that they would cross one or more links (if any) that happen to be monitored by the NSA. In addition, the communications in question must not themselves be blacklisted, as might be the case if they were encrypted, *see* Second Decl. ¶¶ 80, 82, 84, rendering this scenario even more uncertain.

61. Blacklisting by IP address: Mr. Bradner next mentions three scenarios in which he states that blacklisting by IP address would not "guarantee" that the NSA would avoid all interaction with Wikimedia communications. Bradner Reply Decl. ¶ 154(5)(b). The first of these, involving the enclosure of a Wikimedia communication within an MCT, Bradner Reply Decl. ¶ 154(5)(b)(i), is the same as the first of the three scenarios he posited in his first declaration, Bradner Decl. ¶ 367(b)(1), which I have already discussed in my second declaration, Second Decl. ¶¶ 78-80, and again in paragraphs 42-43, above. The second scenario, involving the passage of a Wikimedia communication through an intermediary service such as a VPN, or an email server, Bradner Reply Decl. ¶ 154(5)(b)(ii), is a generic restatement of the second and third scenarios hypothesized in Mr. Bradner's first declaration, Bradner Decl. ¶ 367(b)(2)-(3), which I discussed in my second declaration, Second Decl. ¶¶ 81-84, and again above in paragraphs 42-43. The third blacklisting scenario suggested in Mr. Bradner's second declaration, Bradner Reply Decl. ¶ 154(5)(b(iii), is a repeat of the whitelisting scenario he referred to in paragraph 154(5)(a)(i) of his second declaration, and that I addressed in paragraph 60, above.

62. Port or protocol blacklisting: The last two scenarios hypothesized by Mr. Bradner concern blacklisting by port or protocol (type of communication) rather than by IP address. Bradner Reply Decl. ¶ 154(5)(c). He envisions two situations in which a Wikimedia communication, even if using a blacklisted protocol, would still be copied and scanned during Upstream surveillance if it crossed a monitored link and either (i) it were enclosed in an MCT using a different protocol, one not blacklisted by the NSA, or (ii) it passed through an intermediary (such as an email server) that, likewise, used a different protocol not blacklisted by the NSA.

63. Mr. Bradner gives no examples of either kind of supposed communication, in which a communication using one protocol is transported within another communication using a different protocol. Nor does he give evidence of how frequently such supposed communications could be expected to occur, or under what circumstances. Consequently, the occurrence of such a communication in the first place is itself a matter of speculation. An even greater degree of speculation would then be required to imagine (i) that the protocol used by the enclosing MCT or the intermediary service is not blacklisted by the NSA, (ii) that the IP addresses assigned to the communication within the MCT, or by the intermediary service, are not excluded to due whitelisting, and (iii) that the communication happened to cross an international Internet link monitored by the NSA (if any).

64. For all of the reasons I have explained herein (paragraphs 42-43, 60-63, above), and in my second declaration, Second Decl. ¶¶ 78-85, whether and when any of the scenarios envisioned by Mr. Bradner might come to pass at a particular international Internet link that happened to be monitored by the NSA (if any), such that the NSA would copy and scan communications of Wikimedia's, is a matter of speculation.

CONCLUSION

65. For the reasons I discuss above and in my first two declarations, it remains my opinion that, based on what is publicly known about the NSA's Upstream collection technique, the NSA in theory could be conducting this activity, at least as Wikimedia conceives of it, in a

number of technically feasible, readily implemented ways that could avoid NSA interaction with Wikimedia's online communications.

66. While I offer no opinion on the likelihood that the NSA does or does not, in fact, employ these techniques, I have previously examined, and now re-examined, the bases of Mr. Bradner's opinions (i) that the NSA, in conducting Upstream surveillance, "most likely" copies, reassembles, and scans for selectors all communications packets traversing an international Internet link that is monitored by the NSA (if any); (ii) that it is "implausible" that the NSA uses the traffic-mirroring techniques (white- and blacklisting) described in my first declaration; and (iii) that even if the NSA uses one or more of the techniques I described, it is still "virtually certain" that the NSA copies and scans at least some of Wikimedia's communications. I still conclude that these opinions lack a non-speculative foundation in Internet technology and engineering.

67. My opinions are unaltered by the statements referred to by Mr. Bradner in the FISC's October 2011 opinion and the PCLOB Section 702 Report. As I explained above, the use of traffic-mirroring techniques to implement a filter-then-copy-and-scan approach to Upstream collection would be entirely consistent with both statements. My opinions are unaltered, as well, by the so-called "technical and practical necessities" discussed by Mr. Bradner in his reply declaration. With few exceptions, they are simply reiterations of the same grounds given for the conclusions reached in his second declaration. They, too, are principally based on speculation about the NSA's surveillance practices and priorities, its capabilities and resources, and the number, nature, and communications practices of its Upstream surveillance targets, and lack a non-speculative foundation in Internet technology and engineering.

I declare of penalty of perjury that the foregoing is true and correct to the best of my knowledge and belief.

Executed in New York, New York on March 22, 2019.

HENNING G. SCHULZRINNE