Wikimedia Foundation v. NSA
No. 15-cv-0062-TSE (D. Md.)

Plaintiff's Exhibit 1

IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF MARYLAND

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

REPLY DECLARATION OF SCOTT BRADNER

TABLE OF CONTENTS

1. My name is Scott Bradner. I have been asked by the plaintiff's counsel in *Wikimedia Foundation v. National Security Agency*, No. 1:15-cv-006622-TSE (D. Md.), to provide an expert report addressing the government's reply to the plaintiff's brief and to my declaration, both of which were dated December 18, 2018. My qualifications to express an opinion in the case as well as my compensation and CV are as stated in my previous declaration.

2. A list of the documents provided to me by plaintiff's counsel was attached as Appendix B to the previous declaration.

I. MY CONCLUSION HAS NOT CHANGED.

3. I have carefully reviewed Dr. Schulzrinne's reply declaration as well as the government's reply brief. These documents do not change my basic conclusion in this case that "*it is virtually certain that the NSA has, in the course of the upstream collection program, copied, reassembled and reviewed at least some of Wikimedia's communications*."

II. INTRODUCTION

4. My conclusion rests on four basic foundations:

   i. The NSA is copying packets, reassembling them into communications and then reviewing the communications for the presence of selectors as part of the upstream program.

   ii. Wikimedia's traffic traverses every circuit carrying public Internet traffic into and out of the country (i.e., "international internet links").

   iii. The NSA is monitoring at least one such circuit under the upstream collection program.

   iv. On any circuit it is monitoring, the NSA must be copying, reassembling, and reviewing transactions, including Wikimedia communications, to find those communications that are to or from its targets.

5. The government and Dr. Schulzrinne do not dispute the first two foundations, and they do not seriously dispute the third foundation. But they do dispute the fourth foundation.

6. Dr. Schulzrinne disputes the fourth foundation of my conclusion primarily by describing what I will call a "Wikimedia-avoidance theory"—a hypothetical architecture for an upstream collection program that intentionally avoids Wikimedia's communications (and potentially many other types of communications), rather than having as its goal comprehensively collecting the communications to and from the NSA's targets. This hypothetical architecture is deliberately designed *not* to be comprehensive—because it is designed to avoid entire categories of Internet communications on the off chance that there might be Wikimedia communications present. In offering his Wikimedia-avoidance theory, Dr. Schulzrinne is effectively ignoring the inescapable technical implications of the government's own descriptions of the upstream collection program. Dr. Schulzrinne does not cite any evidence in either of his declarations that the NSA is actually using the extensive filters he describes, nor does he cite any evidence that the NSA is actually avoiding every one of the billions of Wikimedia communications.

7. I disagree with Dr. Schulzrinne and believe that his hypothetical architecture is inconsistent with what the government has disclosed about the upstream collection program. His architecture conflicts with the government's definitive statement that the 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."[1] Dr. Schulzrinne's hypothetical architecture also conflicts with the NSA's goal "*to comprehensively acquire communications that are sent to or from its targets*."[2] The architecture also conflicts with other technical and practical necessities of conducting a program that has collected millions of communications to or from tens of thousands of targets dispersed around the world. Each of these conflicts independently disproves Dr. Schulzrinne's speculation that the NSA is using his Wikimedia-avoidance theory in its upstream collection program, and each independently supports my conclusion concerning the NSA's monitoring of Wikimedia communications.

8. Dr. Schulzrinne argues that his hypothetical architecture, based on extensive whitelist and blacklist filters, does not conflict with government disclosures about the upstream collection program. The bulk of Dr. Schulzrinne's reply declaration is a set of nuanced discussions that are not relevant to the first two conflicts between his hypothetical architecture and the government's disclosures about the upstream collection program. He devotes little space to showing that his hypothetical architecture is consistent with the government's definitive admission that the NSA will acquire wholly domestic communications under some conditions or with the NSA's described goal to comprehensively acquire the communications to or from the NSA's targets. His few arguments relating to these two conflicts are either technically incorrect or consist basically of pleas to ignore the plain meanings of the government's disclosures.

[1] Appendix P (FISC Opinion at 45 (Oct. 3, 2011), *available at* ECF No. 168-4 at 562-643 ("FISC Opinion")).

[2] Appendix F (Privacy and Civil Liberties Oversight Board, *Report on the Surveillance Program Operated Pursuant to Section 702 of FISA* at 10 (July 2, 2014), *available at* ECF No. 168-3 at 199-395 ("PCLOB Report")).

9. As I explain below, Dr. Schulzrinne is also incorrect in his response to the other technical and practical necessities that support my conclusion.

10. Before getting to these explanations, I want to respond to a thread that runs through Dr. Schulzrinne's declaration. He insists that it is impossible for him or me to know the NSA's practices, priorities and capabilities, or the relative likelihood of different technical implementations of the upstream collection program.[3] While absolute assurance may be difficult, the NSA must operate in the real world and deal with the technical and operational limitations inherent in the Internet and in the telecommunications providers it compels to assist it. This need to operate in the real world constrains the ways the NSA could have implemented and be operating the upstream collection program and enables informed deduction of the NSA's actual implementation. Where different implementations present certain technical or practical trade-offs, I have tried to clearly state the degree of certainty or confidence I have in my conclusions.

11. The government has made numerous disclosures relating to the upstream collection program over the years. Many of these disclosures have specific technical implications. The descriptions of the upstream collection program in my first declaration as well as in this declaration are not speculative; instead they are based on the application of my technical expertise to analyze the government's disclosures to understand those technical implications.

[3] Schulzrinne Reply Decl. ¶ 3.

III. DR. SCHULZRINNE'S WIKIMEDIA-AVOIDANCE THEORY CONFLICTS WITH WHAT IS PUBLICLY KNOWN ABOUT THE UPSTREAM COLLECTION PROGRAM.

12. Dr. Schulzrinne's theory of how the NSA might be conducting its upstream collection program conflicts with what the government has publicly said about the upstream collection program. To see why, it is helpful to begin with my original declaration. The final conclusion that I reached in my original declaration was based on four key foundations. I provided support for each of the foundations, summarized here:

A. The four key foundations of my conclusion in my original declaration

1. *Foundation 1: To conduct upstream collection of international public Internet communications traversing a circuit, the NSA must be copying, reassembling, and reviewing transactions on that circuit.*

13. The first foundation for the final conclusion in my previous declaration was "*the NSA must be copying packets, reassembling them into communications and then reviewing the communications for the presence of selectors as part of the upstream program*."

14. I discussed this foundation in my original declaration.[4]

15. In summary, in order to determine if one or more selectors are present in a communication, the NSA must first copy the packets that make up a communication, then reassemble the packets into a copy of the communication. Only after the communication has been reassembled can the NSA review the contents of the communication to determine if the communication contains one or more selectors. The NSA must do this whether or not selectors are present.

[4] Bradner Decl. ¶¶ 250-320, ECF No. 168-2.

16. Neither the government nor Dr. Schulzrinne disputes the requirement that packets be copied and reassembled before they could be reviewed for the presence of selectors.

2. *Foundation 2: Wikimedia communications are transported on all international circuits originating or terminating in the United States.*

17. The second foundation for the final conclusion in my previous declaration was "*Wikimedia's international communications traverse every circuit carrying public Internet traffic on every international cable connecting the U.S. to other countries*" (i.e., "international internet links").

18. I discussed this foundation in my original declaration.[5]

19. In summary, considering the volume of Wikimedia's international communications and the fact that there are users of Wikimedia's U.S.-based services in all of the world's inhabited continents and islands, there must be Wikimedia communications traversing all of the international Internet links connecting the U.S. to the rest of the world.

20. Neither the government nor Dr. Schulzrinne disputes this foundation.

3. *Foundation 3: The NSA conducts upstream collection on at least one international Internet link.*

21. The third foundation for the final conclusion in my previous declaration was "*the NSA is monitoring at least one such circuit under the upstream collection program*."

22. I discussed this foundation in my original declaration.[6]

[5] Bradner Decl. ¶¶ 336-50.

[6] Bradner Decl. ¶¶ 150-153, 222-28, 260-64, 291, 331-35.

23. In summary, based on disclosures the government has made, the NSA is monitoring at least one international Internet circuit that is transporting Wikimedia communications as part of the upstream collection program.

24. Dr. Schulzrinne does not dispute this foundation. The government disputes this foundation and makes, in my opinion, an unpersuasive argument on pages 5 and 6 of their reply brief as to whether government disclosures confirm that the NSA has monitored at least one international Internet circuit.

25. I draw support for my conclusion that the NSA is monitoring at least one such circuit under the upstream collection program from:

a. the NSA's response to plaintiff's Interrogatory No. 12, in which the NSA acknowledges that the "Internet backbone" includes international Internet circuits or links that convey Internet traffic "*internationally via terrestrial or undersea circuits*,"[7]

b. the FISC Opinion of October 3, 2011, in which the FISC stated that "*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*,"[8]

[7] Appendix D (NSA Response to Plaintiff's Interrogatory No. 12, at 18 (Dec. 22, 2017), *available at* ECF No. 168-3 at 77-98).

[8] Appendix P (FISC Opinion at 45).

c. NSA representative Rebecca J. Richards's April 16, 2018 deposition, in which she testified that the "*will acquire*" sentence in the FISC Opinion "*is accurate*,"[9]

d. the PCLOB Report of July 2, 2014, which includes the same "*will acquire*" concession by the government,[10]

e. the NSA's 2014 targeting procedures, which make it clear that, at least in some circumstances, the NSA does not make use of an IP filter to discard wholly domestic communications,[11]

f. the explanation in my first declaration that it is logical and unsurprising that, in designing a program to intercept international Internet communications, the NSA would monitor international Internet links.[12]

26. Taken individually and together, these references make it clear that the NSA has monitored at least one international Internet circuit. Because Wikimedia communications are present on all international Internet circuits, the NSA has monitored at least one international Internet circuit that carries Wikimedia communications.

[9] Appendix K (Transcript of Deposition of Rebecca J. Richards 160:4-17 (Apr. 16, 2018), *available at* ECF No. 168-4 at 105-507 ("Richards Dep.")).

[10] Appendix F (PCLOB Report at 41 n.157).

[11] Appendix T (NSA Section 702 Targeting Procedures at 2 (2014), *available at* ECF No. 168-4 at 1062-1071).

[12] Bradner Decl. ¶¶ 222-24, 293, 332.

4. *Foundation 4: The NSA is copying, reassembling, and reviewing Wikimedia communications on the international Internet links it is monitoring.*

27. The fourth foundation for the final conclusion in my previous declaration was that on any circuit it is monitoring, the NSA must, for a variety of technical reasons, be copying, reassembling, and reviewing all transactions, including Wikimedia communications, to find those communications that are to or from (or about) its targets

28. I discussed this foundation throughout my original declaration.[13]

29. In summary, based on the foundations discussed above, Wikimedia international communications will be transported over at least one of the international Internet circuits the NSA is monitoring. As I explain in detail below, there are several independent reasons why, in the process of monitoring such a link, the NSA must be copying, reassembling, and reviewing at least all international communications transported on the link.

30. Both the government and Dr. Schulzrinne dispute this foundation. Dr. Schulzrinne maintains that the NSA could be using whitelist filters (that is, filters that enumerate the specific IP addresses and/or protocols the NSA wants to review) and/or blacklist filters (that is, filters that enumerate specific IP addresses and/or protocols that the NSA does not want to review) to avoid copying, reassembling and reviewing Wikimedia communications. I disagree that the use of such filters would be technologically consistent with the government's public descriptions of the upstream collection program, and I also disagree that the use of such filters would avoid Wikimedia communications. I will discuss my objections below.

[13] Bradner Decl. ¶¶ 36-48, 272-89, 293-94, 301-18, 333, 335.

31. I will focus on Dr. Schulzrinne's objections, leaving it to counsel to address any of the government's objections that are not based on Dr. Schulzrinne's objections.

B. Why the NSA is copying, reassembling, and reviewing all communications on the international Internet links it is monitoring

32. My conclusion that the NSA is copying, reassembling, and reviewing all communications on at least some of the circuits it is monitoring is supported by at least three independent bases. Each of these bases shows that it is a virtual certainty that the NSA is copying, reassembling, and reviewing Wikimedia's communications. I will now provide a short explanation of each of these bases and of the way in which Dr. Schulzrinne's Wikimedia-avoidance theory conflicts with what is publicly known about the upstream collection program.

1. *The NSA has acknowledged that, on international Internet links it is monitoring, it does not apply IP filters.*

33. The government has disclosed that it does not always apply filters to the traffic on circuits it is monitoring. On at least some circuits, the government has acknowledged that it does not rely on filters to eliminate wholly domestic communications.[14]

34. The government has been quite clear about one of the circumstances in which it does not apply filters. It has acknowledged that it does not rely on filters on the international links it is monitoring.[15]

[14] Bradner Decl. ¶ 291.

[15] Bradner Decl. ¶¶ 292-300.

35. In particular, the NSA has acknowledged to the FISC that it "*will acquire a wholly domestic 'about' communication if the transaction containing the communication is routed through an international Internet link being monitored by NSA*."[16] Note that this disclosure does not say "*may acquire*" or "*might acquire*." Instead, the statement is definitive—the NSA *will* acquire wholly domestic about communications, i.e. wholly domestic communications that include one or more selectors, if they are routed through an international Internet link the NSA is monitoring.[17]

36. Dr. Schulzrinne quotes this statement in ¶ 56 of his declaration. But two paragraphs later he seems to dismiss the categorical nature of the FISC statement and says "*wholly domestic communications of the kinds described above* ***could*** *still be copied and scanned by the NSA*" (emphasis added).[18]

37. There are multiple reasons to believe that the FISC statement should be taken at face value.

38. **Reason for the FISC Opinion**. Considering the circumstances that led to the decision that this quote is part of, it is hard to imagine that the FISC is not being as precise as it possibly could be. The Opinion was the result of a sequence of multiple hearings that were called after the FISC learned of previously undisclosed surveillance activity by the NSA. The hearings involved multiple submissions to the FISC explaining the details of the surveillance activity. It is clear from this that the FISC considered the

[16] Appendix P (FISC Opinion at 45).

[17] Bradner Decl. ¶¶ 292-294.

[18] Schulzrinne Reply Decl. ¶ 58.

situation to be very important and deserving of a very careful Opinion. Thus, there is every reason to believe that the FISC was very careful in what it wrote.

39. **Precise use of language**. It is clear that the FISC was purposeful in its choice of the phrase "*will acquire*." The FISC used a different, less emphatic phrase in a different circumstance only a few paragraphs away.[19]

40. **Richards deposition**. NSA representative Rebecca J. Richards confirmed the sentence in the FISC Opinion was accurate during her April 16, 2018 deposition.[20]

41. The excerpt from the FISC Opinion specifically discusses "about" collection, but that discussion shows that the NSA does not employ IP filters at least in some circumstances.

42. As a technological matter, the only way the NSA *will* acquire wholly domestic "about" communications on the international Internet links it is monitoring, as the government has disclosed, is if it is not applying any filters to that traffic before reviewing the communications to see if they contain one or more selectors. If the NSA were applying, for example, a filter that discarded web (ports 80 and 443) communications, the NSA would miss web-based "about" communication, which would not be consistent with the FISC's statement that all "about" communications "*will*" be acquired. Therefore the NSA must not be applying any filters at these locations. Thus, even in the improbable case where the NSA were deploying the whitelist and blacklist filters of Dr. Schulzrinne's Wikimedia-avoidance theory, based on the FISC Opinion, the filters would not be deployed at the international Internet links being monitored by NSA.

[19] Appendix P (FISC Opinion at 35).

[20] Appendix K (Richards Dep. 160:4-17).

43. The FISC disclosure concerning wholly domestic "about" communications is consistent with other government disclosures that IP filters are not always used, as discussed above. The lack of all filters on these international Internet links means that the NSA would copy, reassemble and review all communications on the link, including all Wikimedia communications that happen to be on the link, for the presence of selectors. This is how the NSA would find the "about" communications. Note that, since it is undisputed that there are Wikimedia communications on every international Internet circuit, there will be Wikimedia communications on any such international Internet circuit that the NSA is monitoring.

44. The definitive statement in the FISC Opinion does not provide any room for any filters, such as the whitelist or blacklist filters Dr. Schulzrinne hypothesizes could be used, because the use of such filters would discard some wholly domestic about communications.

45. Thus, at least on the international Internet links—where it does not employ filters—the NSA must be copying, reassembling and reviewing all communications, including any Wikimedia communications that traverse the link, in order to determine which communications contain targeted selectors.

2. *The PCLOB has explained that upstream collection has been implemented in a technological manner designed to "comprehensively" acquire the communications of the NSA's targets.*

46. The government has said that the aim of the NSA is "*to comprehensively acquire communications that are sent to or from its targets*."[21]

[21] Appendix F (PCLOB Report at 10, 123, 143); Bradner Decl. ¶ 333.

47. Dr. Schulzrinne summarily dismisses the government's own statement as an unrealistic or unrealized goal. He also implies that we should not accept that the PCLOB meant what it said when it used the term "*comprehensively*."[22]

48. I find no reason to doubt the PCLOB's use of the word "*comprehensively*," which it used in describing the technical reasons why the NSA collects "about" communications. Specifically, the PCLOB states that this collection is "*an inevitable byproduct of the government's efforts to comprehensively acquire communications that are sent to or from its targets*."[23] A corollary statement is that the NSA would not be collecting "about" communications if it was not striving to be comprehensive in its collection of the communications of its targets. The PCLOB meant the term "*comprehensively*" to explain the need for the NSA's specific technological implementation of the upstream collection program, and so I find it an appropriate basis on which to explain the technological implementation of the upstream collection program.

49. Even if Dr. Schulzrinne were correct that the NSA is not being comprehensive in which international Internet circuits it is monitoring, it would not follow that the monitoring on the circuits it *does* monitor would not be comprehensive. In fact, the FISC's description of the improper collection of "about" communications only makes sense if the NSA were comprehensively monitoring individual circuits.

[22] Schulzrinne Reply Decl. ¶¶ 72-74.

[23] Appendix F (PCLOB Report at 10).

50. It may be useful to note that the NSA worked "*extensively*" with PCLOB while the PCLOB was preparing the Report to ensure the Report's accuracy.[24] Because of the obvious care that was taken in preparing the Report, it is appropriate to understand the term "*comprehensively*" literally, as the PCLOB obviously intended it to be taken. It is also useful to note that Ms. Richards confirmed in her deposition that the NSA reviewed every sentence of the PCLOB Report, that it identified any inaccuracies in the Report to the PCLOB, and that it does not believe there to be any inaccurate statements of fact in the Report.[25]

51. The use of whitelist and/or blacklist filters as Dr. Schulzrinne hypothesizes is incompatible with a goal of comprehensively acquiring communications that are to or from NSA targets. Such filters work against a goal of being comprehensive. Any such filter inevitably discards communications that might include communications that are to or from the NSA's targets. For example, Dr. Schulzrinne hypothesizes a blacklist filter that would discard traffic to or from Wikimedia servers.[26] He imagines that such a blacklist would discard all communications accessing information on Wikimedia websites by an NSA target. As Dr. Schulzrinne states, I do not know the NSA's surveillance priorities, practices, and capabilities insofar as they are unstated or not inferable based on the NSA's extensive public disclosures.[27] But any such blacklist would, by definition, be incompatible with the government's *stated* goal of completeness.

[24] Appendix K (Richards Dep. 105:20-106:13); Appendix F (PCLOB Report at 3-4).

[25] Appendix K (Richards Dep. 101:22-102:5, 105:7-12, 105:20-106:13, 107:1-5, 108:11-15, 145:9-12).

[26] Schulzrinne Reply Decl. ¶¶ 12-13, 39-42.

[27] Schulzrinne Reply Decl. ¶ 3.

52. The use of whitelist filters, as imagined by Dr. Schulzrinne, is even less compatible with the concept of completeness. At least with a blacklist, one is specifying the relatively few addresses or protocols that are not of interest among the billions of possible addresses and thousands of possible protocols. In that case, communications to or from unknown addresses or using unknown protocols will still be reviewed. With a whitelist, however, one is specifying the relatively few addresses or protocols that are of interest. This means that communications to or from most of the billions of possible addresses or using most of the thousands of possible protocols will be discarded and not reviewed to see if they are from or to NSA targets. It is certainly technically possible to design an Internet surveillance program using a whitelist, but doing so would purposefully ignore most of the Internet, and it would be inconsistent with the publicly known details about the upstream collection program.

53. If the NSA were using a whitelist, or even a blacklist, of the sort that Dr. Schulzrinne speculates, the PCLOB would not have been able to say that the NSA's goal was "*to comprehensively acquire communications that are sent to or from its targets*." I will discuss a number of other flaws in Dr. Schulzrinne's concept of using whitelist and/or blacklist filters below.

54. Based on the technical detail in the PCLOB Report, the NSA must be comprehensively copying, reassembling, and reviewing all communications on a circuit that it is monitoring if the NSA is to meet its goal "*to comprehensively acquire communications that are sent to or from its targets*." This goal is technologically incompatible with any significant use of filters other than those IP filters that ensure that at least one end of a communication is outside the U.S., and, as discussed above, the

government has disclosed that IP filters are not used on international Internet circuits. Since there are Wikimedia communications on every international Internet circuit, there will be Wikimedia communications on any international Internet circuit that the NSA is monitoring and, if the NSA is comprehensively copying, reassembling, and reviewing all packets on circuits that it is monitoring so that it can comprehensively acquire communications that are sent to or from its targets, the NSA will be copying, reassembling, and reviewing Wikimedia communications.

3. *Other technical and practical necessities make clear that the NSA is copying, reassembling, and reviewing Wikimedia's communications.*

55. There is no reliable way for the NSA to know if an individual packet on a circuit is part of a communication that contains one or more selectors without reviewing a reassembled communication containing that packet. Thus, the NSA cannot know in advance which packets need to be copied because they are part of a communication that contains selectors. Thus, the NSA must be copying, reassembling, and reviewing at an absolute minimum those communications it wishes to scan for the presence of selectors.[28]

56. The NSA could not be making use of extensive whitelist and/or blacklist filters, such as the Wikimedia-avoidance architecture Dr. Schulzrinne imagines, because of the technical inability to know in advance which packets on a circuit are part of communications to or from the NSA's targets. Dr. Schulzrinne speculates that the NSA could know its targets' IP addresses and communications protocols in advance, but as I explain at length in the next section, that is not possible given all that we know about the scale and purposes of the upstream collection program.

[28] Bradner Decl. ¶¶ 236-48, 30-18, 333, 335.

57. The only way that the upstream collection program could possibly avoid all of Wikimedia's ubiquitous communications is if the NSA had actively strived to eliminate them, and Dr. Schulzrinne presents no evidence that the NSA has ever attempted to do so or any plausible explanation for why it would do so. Even setting this fact aside, the whitelist and blacklist filters Dr. Schulzrinne imagines would not, in fact, guarantee that the NSA would be able to avoid all Wikimedia communications. Even if they could do so, there is no plausible technical or practical reason why the NSA actually would want to avoid all Wikimedia communications. Some Wikimedia communications, for example, those communications that could reveal what the NSA's foreign intelligence targets are reading and writing on Wikimedia websites such as Wikipedia, would provide information that the NSA could consider to be potentially of interest. In any case, it is almost inconceivable that the NSA went out of its way to try to specifically ensure that upstream collection would not encounter even a single Wikimedia communication.

58. Dr. Schulzrinne's Wikimedia-avoidance architecture is pure speculation, unsupported by any of the government's disclosures concerning the upstream collection program or by any plausible technical or practical consideration. My explanations of the upstream collection program are based on the public record about it and on my expert analysis of that record and of the technology of Internet surveillance. Dr. Schulzrinne's theory, in contrast, is merely a thought experiment, conducted without any consideration of whether the architecture he has imagined would remotely satisfy the purposes of the upstream collection program.

59. I will address Dr. Schulzrinne's responses to this basis for my opinion in much greater detail in the next section.

60. The above three bases show that the NSA is not employing the types of ubiquitous whitelist or blacklist filters that Dr. Schulzrinne imagines.

IV. DR. SCHULZRINNE'S WIKIMEDIA-AVOIDANCE THEORY ASSUMES THE IMPLAUSIBLE USE OF WHITELIST AND BLACKLIST FILTERS.

61. As I explained above (¶¶ 32-60), there are three independent technological bases that support my final conclusion in this case. Dr. Schulzrinne and the government focus most of their attention on the third basis—that technical and practical necessities make clear that the NSA is copying, reassembling, and reviewing Wikimedia's communications—and do not seriously challenge the first two bases. In the previous section, I focused on the first two bases, and in this section I will mostly focus on Dr. Schulzrinne's responses to my third basis. I will also address at least some of the points he attempts to make concerning fine points in the other bases.

62. Most of Dr. Schulzrinne's response to my third basis involves his Wikimedia-avoidance theory for the upstream collection program—an architecture based on whitelist and blacklist filters.

A. Whitelist and blacklist filters

63. Dr. Schulzrinne's declaration largely focuses on the proposition that the NSA could use whitelist and/or blacklist filters to limit the scope of the upstream collection program at least enough to avoid copying, reassembling, and reviewing Wikimedia communications but still be compatible with government disclosures on the operation of the upstream collection program.[29] He also says that the use of such filters would ensure that Wikimedia communications were not among the communications the

[29] See, e.g., Schulzrinne Reply Decl. ¶¶ 50-51.

NSA copies, reassembles and reviews in the process of searching for communications to or from its targets.

64. I disagree on both suggestions. The use of some filters may be compatible with the government disclosures that discuss the use of IP filters to discard wholly domestic communications, but the use of whitelist or blacklist filters is not compatible with the disclosures such as the ones in the FISC Opinion and PCLOB Report that describe monitoring without the use of IP filters.[30] As I explain below, there are other reasons why whitelisting and blacklisting filters are technologically incompatible with what is publicly known about the upstream collection program, and anyway, such filters would not reliably avoid Wikimedia communications.

1. *Filters on international Internet circuits.*

65. Even if it were the case that the NSA was making use of whitelist and/or backlist filters in some circumstances, for example where filters are also used to discard wholly domestic communications, it does not follow that the NSA would add special whitelist or blacklist filters where it is not using IP filters, such as international Internet circuits as noted in the FISC Opinion. (See above at ¶¶ 33-45.)

2. *Whitelist filters are incompatible with the government's public descriptions of the upstream collection program.*

66. As noted above, a whitelist filter is a filter that enumerates the specific IP addresses and/or protocols that the NSA wants to review. All incoming packets that do not have an IP source or destination address that matches an IP address in the filter list, or are using a protocol that is not listed in the filter list of protocols, will be discarded.

[30] Appendix P (FISC Opinion at 45); Appendix F (PCLOB Report at 36-37).

67. There are multiple reasons that the use of whitelist filters would be incompatible with the public descriptions of the NSA's upstream collection program.

a. Whitelist filters are almost useless in the upstream collection program because using them would require that the NSA know unknowable information.

68. As I said in my previous declaration: Whitelisting requires knowing in advance all of the IP addresses that might be used by each of the NSA's targets as well as assuming that those targets are not moving around and thereby changing their IP addresses.

69. Basically, the underlying assumption inherent in the use of whitelist filters is that the NSA has up to date, comprehensive and accurate information on where its targets will be, what sites they will be communicating with and what protocols they will be using in advance of the start of any such communications.[31] This assumption would be impossible to meet for communications to or from the NSA's targets, but even harder to meet for "about" communications because the NSA would have to (1) know which non-targets will be talking about targets and (2) have comprehensive and accurate information on which IP addresses these non-targets will be using, what sites they will be communicating with, and what protocols they will be using, in advance of the start of any such communications. In other words, the use of whitelist filters involves an assumption of precognition.

70. Dr. Schulzrinne seems to think that developing and maintaining whitelists is easy, but it would be virtually impossible to do so for a surveillance program meant to capture the communications of thousands of targets, and in fact impossible to do so for a

[31] Bradner Decl. ¶ 366(d).

program meant to capture the communications of *unknown non-targets* about targets.[32] The examples he provides for an IP address-based whitelist include unspecified individual IP addresses or blocks of IP addresses,[33] the IP addresses of VPN and e-mail servers,[34] and the IP addresses of selected web servers, webmail or chatroom sites.[35] Dr. Schulzrinne's list focuses on servers rather than clients. This makes sense in the same way as looking under a streetlight to find your lost keys makes sense. As long as the NSA's targets use these servers, the NSA will intercept the targets' communications. But there are billions of IP addresses and countless e-mail, web, and other servers in the world, and it is trivial to set up even more of these kinds of servers. Anyone can set up new servers of the kinds Dr. Schulzrinne lists. For example, I have both an e-mail server and a web server in my house. They were easy to set up. If the NSA were restricting itself to the IP addresses of known servers, it would be deliberately foreclosing its ability to capture large amounts of target traffic. Communications making use of new or temporary servers, including those that have been temporarily set up to facilitate communication by NSA targets, would escape the NSA's upstream collection program.

[32] Dr. Schulzrinne speculates that the NSA might obtain "about" communications exchanged by individuals using whitelisted IP addresses (Schulzrinne Reply Decl. ¶ 58), but the PCLOB has made clear that "about" collection permits the NSA to acquire communications between entirely unknown non-targets (Appendix F (PCLOB Report at 121, 126)). It is not possible to whitelist the IP addresses of unknown non-targets in way that would reliably acquire about communications traversing circuits being monitored by the NSA.

[33] Schulzrinne Reply Decl. ¶¶ 8, 43.

[34] Schulzrinne Reply Decl. ¶ 58.

[35] Schulzrinne Reply Decl. ¶¶ 35, 37.

71. Note that Dr. Schulzrinne admits that an IP whitelist that does not include Wikimedia IP addresses would not exclude Wikimedia communications if someone using a whitelisted address communicated with Wikimedia. [36]

72. In theory, protocol-based whitelist filters would not be quite as useless as IP address-based ones for conducting upstream-*style* collection. But, as I pointed out in my previous declaration, there is nothing that restricts Internet users to using assigned port numbers for their applications.[37] Thus, protocol-based whitelist filters could easily miss a lot of communications that the NSA would otherwise want to review, including those using new, non-public or ad hoc protocols—for example, ad hoc protocols used to facilitate the communications of the NSA's targets.

73. That said, the only example Dr. Schulzrinne provides for a protocol whitelist is one for web protocols.[38] From government disclosures, it is already known that the NSA copies, reassembles and reviews web communications, so it is clear that the NSA is not using protocol-based whitelist (or blacklist) filtering to exclude the web protocols.[39] (See ¶ 130 below.)

b. Dr. Schulzrinne's proposed whitelisting is based on other assumptions or simplifications that are also inconsistent with what is known about the upstream collection program.

74. Dr. Schulzrinne makes a number of significant assumptions or simplifications as he argues that the NSA could be using whitelists and blacklists to

[36] Schulzrinne Reply Decl. ¶ 12.

[37] Bradner Decl. ¶ 109.

[38] Schulzrinne Reply Decl. ¶ 10.

[39] Bradner Decl. ¶¶ 314-315.

implement its upstream collection program without copying, reassembling or reviewing Wikimedia communications.

i. *Number of NSA targets.*

75. As I explained in my last declaration, one reason that whitelisting of the sort Dr. Schulzrinne describes is not remotely possible for the upstream collection program has to do with the number and mobility of the NSA's targets.[40] Dr. Schulzrinne responds by implying that the NSA might only have a few targets for the upstream collection program. As he points out, the government has indicated that the NSA has over 120,000 Section 702 targets, but has not stated explicitly that all of these targets are part of the upstream collection program.[41]

76. While that is true, there must be a significant number of upstream collection program targets or they must be prolific communicators—the government has disclosed that 26 million communications were collected under the upstream collection program in 2011.[42] For example, if there were only a thousand upstream collection targets, they would have to average 26,000 communications captured under the upstream collection program each per year in order for the NSA to have collected 26 million communications per year. While the actual number has not been publicly released, based on how many communications were collected in 2011 there are almost certainly tens of thousands of targets for the upstream collection program.

[40] Bradner Decl. ¶ 366(d).

[41] Bradner Decl. ¶ 334; Schulzrinne Reply Decl. ¶ 46.

[42] See, e.g., Appendix F (PCLOB Report at 37); Appendix P (FISC Opinion at 26, 30-34, 73).

ii. *Impossible targeting.*

77. Another reason that whitelisting of the sort Dr. Schulzrinne describes is not remotely possible for the upstream collection program has to do with the requirement that the NSA know, in advance, the IP addresses of its targets or the services the targets are using, even when the targets move around and their IP addresses change.[43] In particular, it would also, as a general rule, be difficult to identify IP addresses that are exclusively used by particular individuals or groups, including NSA targets, which would be required if the aim is to limit the possible copying, reassembly and review to NSA targets.

78. In addition, not all communications with targets will contain a target's IP address. For example, the IP addresses of targets do not appear in the communications when a target is using an intermediary or a communications service that involves multiple hops.[44] Nor do IP addresses associated with targets appear in "about" communications (¶¶ 108-109). In the former case, a whitelist would miss communications to or from the target, and in the latter case, communications about a target. If the whitelist included the IP addresses of intermediary or communications services that involve multiple hops, it would sweep in the communications of all other users of the services, which could include Wikimedia communications.

79. Dr. Schulzrinne argues that the movements of the NSA's targets could be limited to a "*given geographical area*." He notes that the NSA could use a whitelist that

[43] See, e.g., Bradner Decl. ¶¶ 171, 173-74, 229-30, 366(d).

[44] Bradner Decl. ¶¶ 244-7; Appendix P (FISC Opinion at 33-35).

includes a "*set of IP addresses . . . associated with geographical areas where the target is believed to be located*."[45]

80. The use of ranges of IP addresses in a whitelist can simplify the whitelist creation and maintenance, but their use would broaden the range of addresses that would be accepted by the whitelist. Considering the broad geographic distribution of Wikimedia users, the broader the range the more likely that the whitelist would result in the copying and review of Wikimedia communications.

81. Dr. Schulzrinne's suggestion that target movements could be limited to a given geographic area is theoretical at best. There is no particular reason to think that the NSA targets are so limited.

82. But, even if the target's mobility were limited, that does not mean that their use of the Internet would be restricted to any particular range of IP addresses. Because a target could use different ISPs at different times and at different nearby locations and, because the IP addresses that ISPs use are not geographically assigned, a target could move from IP address range to IP address range as they moved around.[46] Moreover, it is reasonable to infer that the NSA's targets are widely distributed across a number of geographic regions, given that the foreign intelligence, counterterrorism, weapons proliferation, and cyber-security uses of Section 702 surveillance implicate foreign governments, organizations, and actors around the world, so that the IP address ranges would be numerous and varied.[47]

---

[45] Schulzrinne Reply Decl. ¶ 47.

[46] Bradner Decl. ¶¶ 159-60.

[47] *'Section 702' Saves Lives, Protects the Nation and Allies*, NSA/Central Security Service (Dec. 12, 2017), https://perma.cc/3JAL-WVV2 ("*'Section 702' Saves Lives*").

83. In general, I do not think it is possible to reliably predict how a user's IP address may change over time or as they move, much less the IP addresses of thousands of users. This is made even more complicated by the fact that the actual users of a given IP address can change over time.

84. Dr. Schulzrinne also suggests that the NSA could be "*whitelisting the IP addresses of websites, webmail services, and/or chatrooms of interest.*" I will note that a webmail service is a website. In addition, many websites and services are now making use of content distribution networks, which, by design, can have different, and changing, IP addresses in different parts of the world. Many websites and services are also making use of cloud-based services, such as Amazon AWS, which also can have multiple and changing IP addresses. Keeping track of the set of IP addresses in use by a particular service at any point in time is, at best, difficult.

iii. *Whitelist complexity and dynamism.*

85. As discussed above in ¶ 84, Dr. Schulzrinne imagines that the NSA could get by with a very simple set of whitelist rules. But he does not mention the fact that the rules would have to be frequently updated as NSA targets were added or removed, or as they changed their locations or methods of operation. He paints a picture of only having to list the IP addresses of some servers along with some IP addresses of a few individuals and ranges of IP addresses.

86. Dr. Schulzrinne implies that the targets of the NSA's upstream collection program do not include individuals.[48] The government's own public documents indicate that the targets include individuals.[49]

[48] Schulzrinne Reply Decl. ¶ 47.

87. Dr. Schulzrinne suggests the NSA could be using whitelists that do not include Wikimedia addresses.[50] A whitelist that included all of the possible IP addresses or IP address ranges that NSA targets could be using yet excluded Wikimedia addresses would be very large indeed since Wikimedia is using only a handful of the approximately 4 billion possible Internet addresses. (There are a bit more than 4 billion possible IP version 4 addresses, as well as billions and billions of times more IP version 6 addresses. For this declaration I will focus on IP version 4 addresses because those are the addresses in most common use.)

88. The only way that a whitelist used in the upstream collection program could be simple while still ensuring that the NSA is comprehensively collecting communications to and from its targets is if the whitelist included most non-U.S. IP addresses and most protocols, in which case there is little reason to have a whitelist in the first place. Of course, any such whitelist would unquestionably include IP addresses and protocols used by Wikimedia users and thus would unquestionably include Wikimedia communications.

c. Whitelist filters assume that the NSA wants to avoid almost all of the world's communications.

89. If the NSA is using a whitelist filter, it means 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. The use of a whitelist in the NSA upstream collection program would be the equivalent of deciding to only look at the material coming through

[49] See, e.g., Office of the Director of National Intelligence, *Section 702 Overview* at 5-7, https://perma.cc/J9X6-YME6 ("*Section 702 Overview*"); *'Section 702' Saves Lives*, *supra* note 47: Appendix F (PCLOB Report at 36); Appendix N (FISC Submission at 4 (Aug. 16, 2011)).
[50] Schulzrinne Reply Decl. ¶ 12.

a few select holes of a sieve. Since whitelisting is specifically designed to ignore most of the Internet, it would be extraordinarily easy for the NSA's targets to avoid being monitored.

3. *Blacklist filters are incompatible with the government's public descriptions of the upstream collection program.*

90. As noted above, a blacklist filter is a filter that enumerates the specific IP addresses and/or protocols that the NSA does not want to review. All incoming packets that have an IP source or destination address that matches an IP address in the filter list or is using a protocol that is listed in the filter list of protocols will be discarded.

91. The government has disclosed that the NSA does impose one type of blacklist filter in at least some circumstances. Wholly domestic communications are filtered out, at least at some—but not all—locations (¶¶ 33-45). This type of filtering can be done with a blacklist that discards packets whose source and destination IP addresses are both within the U.S.

92. There are multiple reasons that any additional use of blacklist filters would be incompatible with the public descriptions of the NSA's upstream collection program and why the use of blacklist filters would not mean that the NSA was avoiding Wikimedia communications. I will now review them.

a. Improbable non-targeting.

93. Dr. Schulzrinne suggests that the NSA could be using blacklist filters to avoid all communications that are to or from the IP addresses of Wikimedia servers.[51] In hypothesizing that the NSA is using such a filter, Dr. Schulzrinne presupposes that the

[51] See, e.g., Schulzrinne Reply Decl. ¶ 12.

NSA would have had a reason to deliberately avoid Wikimedia communications in the upstream collection program.[52] This is different than in a whitelist filter where the NSA is deciding which communications it wants to look at.

94. Dr. Schulzrinne does not provide any evidence, let alone any creditable reason that the NSA would have specifically decided that it did not want to include Wikimedia communications in the upstream collection program. It would be difficult for him to do so because, as he points out, he does not have any specific knowledge of the NSA's priorities. But, based on the available public information about the operation of the upstream collection program and the program's intelligence-gathering purpose, I find it impossible to infer that the NSA would have singled out Wikimedia communications, among the vast array of communications on the Internet, as communications that should be ignored.

95. Dr. Schulzrinne does talk about reducing the load on the devices the NSA is using in the upstream collection program, but he does not indicate why a desire to reduce load would have led the NSA to exclude Wikimedia communications.[53]

96. Dr. Schulzrinne notes that the inter-regional Internet capacity is very large these days—as much as 295 terabits per second.[54] To put the volume of Wikimedia communications into context, in the six-month period between August 1, 2017 and January 31, 2018, Wikimedia engaged in approximately 760 billion international communications.[55] This works out to about 48 thousand communications per second.

[52] Bradner Decl. ¶ 367(a).

[53] Schulzrinne Reply Decl. ¶¶ 20-22.

[54] Schulzrinne Reply Decl. ¶ 20.

[55] Bradner Decl. ¶ 346.

Assuming that an average Wikimedia communication is 81 packets,[56] and that the average packet length is 1,500 8-bit bytes, that means Wikimedia communication averages about 47 Gbps per second. Thus, Wikimedia represents about 0.016% of the inter-regional Internet capacity. Even if I were way off in my estimate of an average communication length of 81 packets, the percentage of inter-regional Internet capacity, measured in bits per second, represented by Wikimedia communications is still extraordinarily low. Blacklisting the Wikimedia IP addresses would not make any measurable difference to the load experienced by the NSA's upstream collection system. Thus, it is very unlikely that the NSA would have decided to specifically blacklist Wikimedia communications to reduce the load on the upstream collection program systems, even if, as Dr. Schulzrinne suggests, it might be easy to do.[57] Note also that although Dr. Schulzrinne says that, according to a statistics website, one of the Wikimedia websites is the 5th most popular website globally, it does not follow that Wikipedia generates a significant amount of traffic measured in bits per second compared to the inter-regional Internet capacity.

b. Blacklist filters would not guarantee that the NSA would avoid copying, reassembling, and reviewing Wikimedia communications.

97. I noted a number of situations under which NSA's use of a blacklist filter designed to block the communications to or from the IP addresses of Wikimedia servers would not guarantee that the NSA would avoid copying, reassembling and reviewing

[56] Bradner Decl. ¶¶ 144-45.

[57] Schulzrinne Reply Decl. ¶ 41.

Wikimedia communications in my previous report.[58] Dr. Schulzrinne responded to the discussion in my previous declaration in his reply declaration.[59] In his response, Dr. Schulzrinne addressed three of the cases in which I said that Wikimedia communications could still be copied, reassembled and reviewed even if there were a blacklist filter in place that discarded communications to or from the IP addresses of Wikimedia servers. Dr. Schulzrinne responds to my MCT, e-mail and VPN examples.

98. For each of these three cases, Dr. Schulzrinne creates a list of conditions that he says must be true for Wikimedia communications to be copied, reassembled and reviewed.

99. For example, in each of his sets of conditions, he says that the relevant communication would have to *not* be blacklisted in order for that communication to be acquired, as though that would be a difficult condition to meet. But, in reality, the possibility that these communications *would* be blacklisted in each example is far-fetched. For the MCT example, in which a Wikimedia communication is enclosed in an MCT traversing an international Internet link, Dr. Schulzrinne says that the MCT would have to not be blacklisted. Such an MCT could be one to or from a mail server. In the e-mail example, Dr. Schulzrinne says that the e-mail itself would have to not be blacklisted. In this case the communication would also be to or from an e-mail server. And, in the VPN example, he says that the VPN server communications would have to not be blacklisted.

[58] Bradner Decl. ¶ 367(b).

[59] Schulzrinne Reply Decl. ¶¶ 78-87.

100. It is technically possible for the NSA to block communications to and from e-mail servers and to or from VPN servers but it is hard to understand why the NSA would do this. These types of servers are exactly the types of services that NSA targets could be using, and so a world-wide blacklist blocking communications to and from e-mail and VPN servers would block just the type of sites bad actors would be using and the NSA would have an incentive to target.

101. All three of Dr. Schulzrinne's sets of conditions include the obvious requirements that either the user or the server be located outside the U.S., but not both, that communications between the user and server transit an international Internet link the NSA was monitoring, and that the users must be communicating with Wikimedia. These are all requirements I assumed when I described the cases. In my opinion, these requirements are obvious and would likely be frequently met.

4. *Still copying packets.*

102. As I discuss below and in my original declaration, even if the NSA were employing some sort of filter that discarded packets that were part of Wikimedia communications, the NSA would most likely be copying those packets before discarding them.[60]

5. *Probing for blind spots.*

103. In my original declaration I mentioned that if the NSA were using blacklist or whitelist filters to ignore protocols or IP address ranges, targets could probe to see if they could discover the lapses in coverage.[61] Dr. Schulzrinne questions the

[60] See ¶¶ 114-122; Bradner Decl. ¶¶ 269-279.

[61] Bradner Decl. ¶ 366(b).

possibility that the NSA's targets are sophisticated enough to probe for gaps in the NSA's coverage.[62] I will note that foreign intelligence officers and services are among the people and groups that are lawful targets for the upstream collection program,[63] and they are among our most sophisticated adversaries. I will also note that such probing is more likely to require process rather than technical sophistication. For example, it could involve purposefully communicating information about an information resource over a protocol you suspect the NSA is not monitoring and then monitoring the information resource to see if it is accessed, or purposefully conveying actionable information, such as the identity of a foreign agent, over such a protocol and seeing if action is taken against the agent.

104. There is another significant risk to the NSA's use of whitelists and/or blacklists to limit what its surveillance devices copy, reassemble, and review, and it would not require probing.

105. To implement the kind of whitelisting and blacklisting that Dr. Schulzrinne hypothesizes in the way that he hypothesizes, ISPs would need to configure their routers with the whitelists and blacklists. It is standard practice for ISPs to backup their router configurations in their network management systems, so that they can quickly deploy, modify or re-deploy the configurations as needed. There have been too many cases where ISPs' network management systems have been compromised, and so whitelisting and blacklisting of the sort Dr. Schulzrinne describes would create the unnecessary risk of compromise of the NSA's whitelists and blacklists. In addition,

[62] Schulzrinne Reply Decl. ¶ 32.

[63] 50 U.S.C. § 1881a.

multiple ISP technicians generally have access to the management stations, increasing the number of people that would have to be trusted.

106. Dr. Schulzrinne strangely suggests that the NSA might not care if foreign actors knew of ways to get around NSA monitoring.[64] I agree that I am not privy to the NSA's priorities that are not public, but I would find it quite extraordinary if the NSA did not care whether targets of its surveillance had a roadmap for evading its surveillance.

107. In fact, George C. Barnes, Deputy Director of the NSA, specifically stated that revealing information that could be used to help an adversary evade the NSA would be a problem:

> *Revealing which channels* [of communication] *are free from surveillance and which are not could also reveal sensitive intelligence methods, and thereby help an adversary evade detection and capitalize on limitations in the NSA's surveillance capabilities*.[65]

B. "About" communications

108. Dr. Schulzrinne's proposed use of whitelist and blacklist filters is also entirely inconsistent with "about" collection. About collection is a process within the upstream collection program that involves the collection of communications between two non-targets that contain one or more selectors associated with an NSA target. Dr. Schulzrinne discusses "about" communications in ¶¶ 49-52 of his reply declaration.

109. As a general rule, the IP addresses on the packets that make up the "about" communications will likely have no relation to any targets. For this reason, the use of

[64] Schulzrinne Reply Decl. ¶ 33.

[65] Barnes Decl. ¶ 57, ECF No. 141-1.

whitelist and blacklist filters of the sort that Dr. Schulzrinne describes is not compatible with "about" collection, because that kind of filtering would guarantee that the NSA's upstream collection devices miss "about" communications.

110. Dr. Schulzrinne responds to this fact by describing a two-step process for the collection of "about" communications that he claims is compatible with the use of whitelist and blacklist filters. The first step uses a whitelist IP address-based filter, but to set up this whitelist filter the NSA would have to know in advance which non-targets' IP addresses to whitelist, or what servers' IP addresses to whitelist in order to find the "about" communications.

111. Dr. Schulzrinne oversimplifies the problem by saying that "about" communications would be collected if the communication containing the "about" selector were whitelisted. Dr. Schulzrinne ignores the fact that in order to include such communications in a whitelist, the NSA would first have to know in advance which non-targets were going to be talking about targets and also know in advance what IP addresses the non-targets would be using. If the NSA were following Dr. Schulzrinne's description, they might capture an occasional "about" communication if one of the non-targets was using an IP address or service that the NSA had whitelisted for some other reason, but they could not normally capture them.

112. Under Dr. Schulzrinne's model, the only way the NSA could reliably capture about communications would be to whitelist all non-wholly domestic communications, in which case the whitelist would be guaranteed to result in the copying and review of Wikimedia communications.

V. DR. SCHULZRINNE'S WIKIMEDIA-AVOIDANCE THEORY IS IMPLAUSIBLE FOR NUMEROUS OTHER TECHNICAL AND PRACTICAL REASONS.

113. Dr. Schulzrinne made a number of other points in his reply brief. I will respond to some of them now.

A. Copy-then-filter vs. in-line filter

114. In my original declaration, I expressed the opinion that the NSA was most likely using a copy-then-filter architecture for its upstream collection program.[66] The government treats this as a significant concession, but the government completely misrepresents how this point relates to my ultimate conclusion. My discussion of the copy-then-filter implementation is a discussion of an independent reason to believe that the NSA is copying Wikimedia's communications as they travel across international Internet links—even if the NSA were performing filtering of the kind Dr. Schulzrinne hypothesizes. Specifically, a copy-then-filter architecture renders all of Dr. Schulzrinne's speculation about filtering irrelevant. That is because a copy-then-filter architecture involves the NSA copying all communications on a circuit independent of any filtering that might subsequently be performed on the copies of the packets. That of course includes the copying of Wikimedia's communications.

115. It bears emphasis, however, that even if the NSA is not using a copy-then-filter architecture, but instead is using the in-line filtering architecture Dr. Schulzrinne describes (which he refers to as "*filter-then-copy-and-scan*"), some Wikimedia communications will be copied. Dr. Schulzrinne's Wikimedia-avoidance architecture is entirely implausible for the many independent reasons I have explained above (including

[66] Bradner Decl. ¶¶ 265-289.

the multiple ways in which his theory conflicts with the government's public disclosures) and explain below (including the multiple ways in which his theories are at odds with the technical and practical necessities of conducting the upstream collection program as it has been described). In short, even if implemented as Dr. Schulzrinne imagines, Dr. Schulzrinne's filters will miss some Wikimedia communications and those communications will be copied, reassembled and reviewed.

116. I will now respond to Dr. Schulzrinne's points about the implications of using the one implementation (copy-then-filter) versus the other (in-line filtering). But, it continues to be my opinion that the copy-then-filter architecture is the simplest, most reliable and easiest to operate architecture for the NSA to use for the upstream collection program. If the NSA uses this architecture, all packets on a communications circuit being monitored by the NSA, including the packets that make up Wikimedia communications, are copied. It is possible that the NSA uses different architectures in different monitoring locations, but, in my opinion, the advantages of the copy-then-filter architecture mean that it is most likely the default architecture. Even if the copy-then-filter architecture is not being used everywhere, all packets on the circuits where the copy-then-filter architecture is used are copied, as are many packets on the circuits where an in-line filter architecture is used.

1. *Fiber-optic splitter.*

117. The simplest option involves the use of a fiber-optic splitter, which Dr. Schulzrinne suggested in his first declaration and I discussed in my original declaration.[67] A fiber-optic splitter produces a copy of all communications on a fiber by splitting the

[67] Schulzrinne Decl. ¶¶ 55-56, ECF No. 164-4; Bradner Decl. ¶¶ 275-77.

light on the fiber into two streams.[68] Dr. Schulzrinne notes that adding a fiber-optic splitter into a network adds a potential failure point and causes loss of optical power.[69] He is correct that a fiber-optic splitter could be a failure point, but a fiber-optic splitter is a passive device that does not include a processor or software that could have bugs or need upgrading, and it does not require any configuration or power so the probability of failure is very low and the possibility that misconfiguration or that a power failure could impact the network is nonexistent. A fiber-optic splitter does reduce the optical power that would be received by the ISP's receiving device but, as long as the ISP knows the splitter is in-line, the receiving device can be configured to compensate for the loss.

118. On the other hand, using a router's mirror function (as Dr. Schulzrinne describes), would have a much higher failure and disruption probability because the router requires power and because the router includes a computer and software that can have bugs, would need updating, and would be vulnerable to hacking. If the ISP used an existing router to filter and then copy communications, the *added* risks would be a little bit less significant, but the need to constantly reconfigure the device with updated blacklists and whitelists would create the risk of misconfiguration or overloading. Either way, the risks of failure are greater for the in-line device Dr. Schulzrinne proposes than for a fiber-optic splitter.

119. Dr. Schulzrinne notes, as I did in my original declaration, that the use of a fiber-optic splitter would mean that it would need to be coupled with an opto-electronic

[68] Bradner Decl. ¶¶ 5, 275-77.

[69] Schulzrinne Reply Decl. ¶ 27.

device to split out the channels the NSA wanted to monitor.[70] Dr. Schulzrinne paints a pretty dire picture of the requirements for this device, but, as I pointed out in my original declaration, not all channels on the cables are used to transport international Internet communications.[71] The splitting device only needs to pay attention to the circuits that are so used and, of those, only the circuits that the NSA wishes to monitor. Note that the splitting device Dr. Schulzrinne mentions is not an esoteric piece of equipment; it is the same device that ISPs routinely use to split the light on optical fiber into different channels and is normally included in the router that the optical fiber is plugged into.

2. *Configuring the ISP router to mirror communications.*

120. The other option to support the copy-then-filter architecture is, as Dr. Schulzrinne suggests, for the NSA to command the ISP to configure its router to mirror the communications on the circuits the NSA wants to monitor and send the mirrored packets to one or more NSA-operated devices.[72]

121. The operationally simplest way to do this is to command the ISP to configure its router to mirror all of the packets on the channels the NSA wants to monitor. Such a configuration is very simple and would not have to change over time. This way also means that the ISP personnel are not exposed to any of the NSA's collection methods other than the fact that data collection is being done on a particular circuit.

[70] Bradner Decl. ¶ 277; Schulzrinne Reply Decl. ¶ 21.

[71] Bradner Decl. ¶¶ 214-25.

[72] Bradner Decl. ¶ 278.

3. *All packets are copied.*

122. All packets on a circuit, including the packets comprising any Wikimedia communications, are copied with either of the above two designs.

4. *Configuring the ISP router to filter before mirroring.*

123. Dr. Schulzrinne hypothesizes that the NSA could be applying whitelist or blacklist filters to the packets *before* the router mirrors the packets to the NSA upstream collection devices.[73] He says, and I agree, that if the NSA is doing so, then the logical place to apply such filters is in an ISP router that will be processing the stream of packets entering or exiting a communications channel the NSA wants to monitor. Dr. Schulzrinne and I disagree as to whether the NSA could in fact be using such whitelist or blacklist filters on all of the circuits that it monitors, even if it uses them on some, and we disagree as to what specific configurations would be likely for whitelist or blacklist filters. I discuss this disagreement at length above.[74]

124. There are different categories of whitelist or blacklist filters that have different implications when it comes to configuring the filters. In this case, I believe it most likely that the actual router configuration is being performed by ISP personnel since, in my experience, an ISP would be unlikely to allow non-ISP personnel to configure its routers. In theory, the NSA could create some configuration data files that the ISP personnel could then load into the router—this would put the NSA directly in charge of the filter details while avoiding having the fingers of non-ISP personnel in the routers. In this case the ISP personnel would still have access to the details of the

[73] Schulzrinne Reply Decl. ¶ 28.

[74] See ¶¶ 63-107.

configuration. In any case, the more complex and detailed the filter configurations, the more often they will change. If an ISP had to implement all of the whitelist and blacklist filters Dr. Schulzrinne has suggested, the ISP would have to be updating its router configuration all the time. Operationally, this is not a good idea from an ISP's perspective because it increases the chance of human error impacting the ISP operations.

125. In his reply declaration, Dr. Schulzrinne says that he does not suggest installing an NSA-operated device in the middle of an ISP's network.[75] There are only two options for the type of in-line filter that Dr. Schulzrinne proposes: (1) having an NSA-operated device perform the filtering, or (2) exposing ISP personnel to the details of the NSA's collection program. Each of these options has significant downsides. If in-line filters are used at all, it may well be that there is no single answer. For example, the NSA might want to operate its own device in smaller ISPs or ISPs demonstrating less technical expertise but delegate the operation to ISP personnel in larger ISPs or those with more technical expertise.

126. As I discussed in my previous declaration, it is my opinion that the NSA would want to limit the exposure of the details of at least some of the types of whitelist and blacklist filters Dr. Schulzrinne suggests could be used. Dr. Schulzrinne dismisses that concern by pointing to the fact that the NSA shares target information with telecommunications providers under the PRISM program.[76] But Dr. Schulzrinne ignores the fact that the categories of information are quite different. With PRISM, the identities of one or more individual targets are exposed to the personnel of the telecommunications

[75] Schulzrinne Reply Decl. ¶ 23.

[76] Schulzrinne Reply Decl. ¶ 18.

companies it compels to provide assistance. But the detailed configuration of blacklist and whitelist filters can provide a roadmap that can be used to entirely avoid NSA surveillance.

127. Note that the NSA has to share confidential target information with telecommunications providers under PRISM, since only the telecommunications providers have the ability to retrieve the communications of interest. The NSA does not have that constraint in the upstream collection program. The NSA could:

a. use a copy-then-filter approach with a fiber-optic splitter, which would reveal no confidential information to the provider other than the fact that the NSA was monitoring one or more channels on a fiber;

b. use a copy-then-filter approach by commanding the ISP to provide the NSA with a full copy of the packets on a particular channel, which would increase the shared confidential information to include the specific channels being monitored;

c. operate its own in-line filter device which would also share the information about which channels were being monitored; or

d. command the ISP to operate an in-line filter, which would require the NSA to share the details of exactly what it is monitoring and not monitoring (i.e., the configurations for the whitelist and blacklist filters Dr. Schulzrinne hypothesizes).

128. Given the fact that the NSA has a choice, it seems reasonable to infer that the NSA would want to minimize the confidential information that it had to share to

operate the upstream collection program. (See also ¶ 107 above, citing Barnes Decl. ¶ 57.)

129. It is one thing to expose the fact that the NSA has asked for Mr. Smith's e-mail and a very different thing to expose the fact that the NSA is using a whitelist filter that discards packets to or from large chunks of the Internet or discards all packets that are not e-mail or web traffic.

B. Collecting "web activity"

130. Dr. Schulzrinne questions the idea that the NSA might be monitoring communications to or from web servers.[77] I discussed the government's disclosure that it is monitoring "web activity" in my original declaration.[78] Dr. Schulzrinne suggests that the government might have been sloppy and was referring to overall Internet activity when they wrote "web activity."[79] That seems to me to be a very tenuous argument. I have seen no indication in any of the government's released documents that they are that sloppy. In particular, the reference to "web activity" is in a formal and highly technical government submission to the FISC in response to a highly technical request from the court, hardly a place that anyone would be sloppy. And, the context in which the phrase was used makes total sense if the government was using the phrase precisely, as a reference to the world wide web protocols (HTTP/HTTPS).

131. It is also well known that terrorists make use of communications tools that use HTTP/HTTPS. Most human-to-human communications on the Internet are transported using HTTP/HTTPS, whether over older mediums (websites) or more recent

[77] Schulzrinne Reply Decl. ¶ 36(b).

[78] Bradner Decl. ¶¶ 314-15.

[79] Schulzrinne Reply Decl. ¶ 36(b).

ones (messaging services). For example, terrorists make use of the Telegram application, which can operate over port 80 (HTTP).[80] There are obviously examples of widely used protocols that do not use HTTP/HTTPS—for example, FTP/SMTP/IMAP/POP/etc. — but an increasing amount of Internet communications do use HTTP/HTTPS to increase security (HTTPS) or to bypass firewalls (both).

C. Collecting web communications

132. Dr. Schulzrinne spends some time discussing ways the NSA could be filtering out web or encrypted traffic to make it less likely that it is copying, reassembling and reviewing Wikimedia communications. I will now review his suggestions.

133. Dr. Schulzrinne suggests that the NSA might be configuring a blacklist to block both HTTP and HTTPS communications.[81] First of all, the government has acknowledged that it is capturing web activity, i.e. HTTP and/or HTTPS communications (¶ 130). The NSA cannot be both blocking all web communications and be collecting web activity.

134. Note that web traffic makes up a very large percentage of Internet traffic. For example, one study of traffic between the Japanese WIDE Project ISP and its backbone ISP showed that web traffic was 75% of the overall traffic.[82] As Dr. Schulzrinne notes, web traffic is not restricted to web sites such as cnn.com and

[80] Joby Warrick, *The 'App of Choice' for Jihadists: ISIS Seizes on Internet Tool to Promote Terror*, Wash. Post (Dec. 23, 2016), http://wapo.st/2hzoY6P; *MTProto Mobile Protocol: Transport*, Telegram, https://perma.cc/T6FL-WWP8.

[81] Schulzrinne Reply Decl. at 15 n.7.

[82] Chia-ling Chan, et al., *Monitoring TLS Adoption Using Backbone and Edge Traffic* (2018), https://perma.cc/6C8D-GWCT.

wikipedia.org; web protocols are also used for webmail and chatrooms.[83] Ignoring web traffic as Dr. Schulzrinne has suggested would exclude the vast quantities of human-to-human communications that are transported by the web.

135. As I mentioned in my original declaration, any such discarding of all web communications would leave a very large hole in the NSA's coverage[84]—contrary to any notion of completeness such as that noted in the PCLOB Report

136. Dr. Schulzrinne suggests that the NSA could be restricting its collection of web communications to a few sites such as "*specific webmail and chatroom sites*." As I discussed above, this would also leave very large holes in the NSA's coverage. (¶ 134) Any such filtering would also be contrary to the aim of the "about" collection program, which is to collect communications between unknown non-targets.

D. Collecting encrypted communications

137. Dr. Schulzrinne says that the NSA might not be collecting HTTPS communications using the authority in Section 702 to collect encrypted communications.[85] Dr. Schulzrinne makes no actual argument that the NSA is not collecting at least some encrypted communications, he just says that my arguments are not technical.

138. Multiple government disclosures make it clear that the NSA collects encrypted communications. For example:

a. The PCLOB noted: "*With respect to all of the agencies, extensions from these age-off requirements may be sought from*

[83] Schulzrinne Reply Decl. ¶ 35.

[84] Bradner Decl. ¶ 366(f).

[85] Schulzrinne Reply Decl. ¶ 36(a).

a high-level agency official. Other limited exceptions apply, such as to communications that are still being decrypted." (Note the phrase "*still being decrypted.*")[86]

b. The PCLOB also noted: "*The NSA may also retain communications beyond the normal age-off period if it is still decrypting the communication or using the communication to decrypt other communications*." (Note the phrases "*still decrypting*" and "*decrypt other communications.*")[87]

c. The NSA's minimization procedures note: "*In the context of a cryptanalytic effort, maintenance of technical data bases requires retention of all communications that are enciphered or reasonably believed to contain secret meaning, and sufficient duration may consist of any period of time during which encrypted material is subject to, or of use in, cryptanalysis*." (Note the phrase "*encrypted material is subject to*.")[88]

139. Dr. Schulzrinne hypotheses that while the NSA may collect encrypted communications it may only be doing so under PRISM.[89] This seems unlikely. The telecommunications providers assisting the NSA in the case of PRISM will frequently have direct access to the user's unencrypted communications, for example in an e-mail

[86] Appendix F (PCLOB Report at 60).

[87] Appendix F (PCLOB Report at 63).

[88] Appendix S (NSA Section 702 Minimization Procedures at 10 (2014), *available at* ECF No. 168-4 at 1046-1061).

[89] Schulzrinne Reply Decl. ¶ 36(a).

server, whereas communications across the Internet are increasingly being encrypted.[90] If the upstream collection program were to ignore encrypted communications it would be increasingly unable to collect any communications.

E. GCHQ surveillance

140. I discussed some public disclosures from the U.K.'s signals intelligence agency, Government Communications Headquarters (GCHQ), in my first declaration in order to "*reinforce my conclusions that the NSA relies on the copy-then-filter configuration to conduct the upstream collection program and that it does not selectively filter traffic prior to copying it as Dr. Schulzrinne hypothesizes it could*."[91]

141. Those include the disclosure by the GCHQ that, under a surveillance program analogous to upstream collection, "*it is necessary to intercept the entire contents of a bearer [circuit], in order to extract even a single specific communication for examination*."[92]

142. Dr. Schulzrinne dismisses the disclosures as only being "*the roughest outline*" of the process the GCHQ uses, and as being "*non-technical*."[93] I disagree and find the disclosures have enough detail for me to draw my conclusions.

143. He also says that the disclosures, even if not detailed, are "*quite comparable*" to his suggested filter first approach. I disagree with this as well.

144. Dr. Schulzrinne provides a citation describing the GCHQ's bulk interception as purported proof of his contention.[94] The citation Dr. Schulzrinne provides

[90] John Maddison, *Encrypted Traffic Reaches a New Threshold*, NETWORKComputing (Nov. 28, 2018), https://perma.cc/6VFY-YEGL.

[91] Bradner Decl. ¶ 369.

[92] Bradner Decl. ¶ 368.

[93] Schulzrinne Reply Decl. ¶¶ 59-60.

is from a filing by the U.K. government in the European Court of Human Rights. The U.K. filing, in turn, refers to a "*Bulk Powers Review*" of "*the operational case for various intelligence gathering powers.*"[95] The U.K. filing quotes the *Bulk Powers Review* to provide a summary on how the GCHQ interception program works.[96] Dr. Schulzrinne's citation is of the U.K. filing's quoting of the *Bulk Powers Review*.

145. Dr. Schulzrinne says that the U.K. filing "*actually describe*[*s*] *a collection approach quite comparable (at least at a general level) to the type of IP address and port and protocol number filtering described in my earlier declaration*."[97] I disagree with Dr. Schulzrinne's analysis of the U.K. filing he provides in support of his conclusion. I do not believe that the citation shows any evidence that the GCHQ is filtering traffic on a channel before copying the traffic; in fact, the citation shows the opposite. The paragraph in the *Bulk Powers Review* that immediately follows the outline Dr. Schulzrinne cites makes this clear:

> ***The two major processes***
>
> *2.19. A description is given in the 2015 ISC report (paras 61-73), of two major and distinct processes that apply to interception under bulk warrants. Those processes are identified in more detail in the closed version of the report, and I have been briefed on each of them. In summary:*
>
> *(a)* ***The "strong selector" process*** *(2015 ISC report, paras 61-64) operates on the bearers that GCHQ has chosen to access. As the*

---

[94] Schulzrinne Reply Decl. ¶ 61.

[95] David Anderson Q.C., U.K. Independent Reviewer of Terrorism Legislation, *Report of the Bulk Powers Review*, Cm 9326 (August 2016), https://perma.cc/V3ME-QZED ("*Bulk Powers Review*").

[96] *Bulk Powers Review*, *supra* note 101, at 23-24.

[97] Schulzrinne Reply Decl. ¶ 61.

internet traffic flows along those chosen bearers, the system compares the communications against a list of strong selectors in near real-time. Any communications which match the selectors are automatically collected and all other communications are automatically discarded. The nature of the global internet means that the route a particular communication will take cannot be predicted and a single communication is broken down into packets which can take different routes. In order to identify and reconstruct the wanted communications of subjects of intelligence interest, GCHQ's processing relies on accessing the "related communications data" (secondary data) in the bearer.

A copy of all the communications on a bearer has to be held for a short period in order to allow the strong selectors to be applied to those communications. This process accordingly requires a bulk warrant under the Bill. However, in the opinion of the ISC, "while this process has been described as bulk interception because of the numbers of communications it covers, it is nevertheless targeted since the selectors used relate to individual targets".

(b) ***The "complex query" process*** *(2015 ISC report paras 65-73) is used where GCHQ is looking to match much more complicated criteria, for example with three or four elements. This process operates across a far smaller number of bearers. These bearers are not chosen at random, as GCHQ focuses its resources on those most likely to carry communications of intelligence value. As a first step in the processing under this method the system applies an initial set of processing rules. Those rules seek to select communications of potential intelligence value while discarding those least likely to be of intelligence value. The selected communications are not available to GCHQ staff to search through at will. Further complex searches draw out the*

> *communications of intelligence value. By performing searches combining a number of criteria, the odds of a 'false positive' are considerably reduced.*
>
> *This second process is closer to true bulk interception, since it involves the collection of unselected content and/or secondary data. It permits types of analysis and selection that are not currently achievable in the near real-time environment of the strong selector process (2.19(a) above). But as with the first process, it remains the case that communications unlikely to be of intelligence value are discarded as soon as that becomes apparent.*[98]

146. The description of the "strong selector" process specifically says "*[a] copy of all the communications on a bearer has to be held for a short period in order to allow the strong selectors to be applied to those communications."* Combining this statement with the description of the third stage of collection in the extract Dr. Schulzrinne provided makes it clear that all of the communications on a bearer (GCHQ's term for a circuit) are copied and stored at least temporarily so that those communications that contain selectors, if any, can be located.

147. The description of the "complex query" process says that the process is "*closer to true bulk interception, since it involves the collection of unselected content and/or secondary data.*"[99] This statement by itself notes that the GCHQ is collecting "*unselected content*." To do so, it is copying communications that it has not checked for the presence of selectors.

[98] *Bulk Powers Review*, *supra* note 101, at 24-25.

[99] *Bulk Powers Review*, *supra* note 101, at 25.

148. In both cases, GCHQ is copying all the contents of a bearer, which is the point I made in my original declaration.

F. ISP-operated copy device

149. Dr. Schulzrinne suggests that a fiber-optic splitter, or even an electronic device such as a router, could be operated by the ISP and the copy of the communications created by such a device could be sent to a filtering device operated by the ISP with the output of the filtering device sent to the NSA.[100] Such an arrangement would not actually change the fact that the NSA is creating a copy, since the copy device and filter would be operated at the direction and auspices of the NSA. As I noted in my first declaration, work performed at the direction of the NSA is still work done by the NSA.[101] In addition, as with other copy-then-filter configurations, all the packets on the circuit, including packets that are part of Wikimedia communications, are copied.

G. EINSTEIN

150. I mentioned in passing the U.S. government-operated EINSTEIN 2 & 3 systems in my original declaration.[102] Dr. Schulzrinne made rather much more of the mention than I had in mind. I just mentioned EINSTEIN as an example of a deep packet inspection (DPI) device. But I will comment on Dr. Schulzrinne's discussion of EINSTEIN.[103]

151. Dr. Schulzrinne tries to differentiate EINSTEIN from the NSA's upstream collection program in two ways: (1) he says that EINSTEIN has to look at all traffic

[100] Schulzrinne Reply Decl. ¶ 64.

[101] Bradner Decl. ¶ 5.

[102] Bradner Decl. ¶¶ 259, 286.

[103] Schulzrinne Reply Decl. ¶¶ 67-69.

whereas the upstream collection program, at least as he imagines it, does not; and (2) EINSTEIN is a cybersecurity system and upstream collection program is not.

152. Relating to (1), Dr. Schulzrinne says "*cyber attacks can use any protocol, originate from any external Internet host, and can target any destination system, to be effective an intrusion-detection system must inspect all incoming traffic*."[104] Of course, the same can be said, with the exception of the discarding of wholly domestic communications, of the upstream collection program. The main theme of Dr. Schulzrinne's declaration is how to limit the upstream collection program so that it avoids Wikimedia communications at the expense of being able to capture communications from, to or about its targets which "*can use any protocol, originate from any external Internet host, and can target any destination system*."[105]

153. Relating to (2), the use of DPI is not limited to cybersecurity systems. The upstream collection program uses DPI to find the selectors in the communications it reviews. In addition, the NSA's public materials include mention of the use of the upstream collection program for cybersecurity purposes.[106]

VI. SUMMARY

154. In order for Dr. Schulzrinne to be correct in his speculation that the NSA could operate the upstream collection program without ever copying, reassembling, or reviewing even a single Wikimedia communication, every one of the following assumptions, which are prerequisites for his claim, must be true:

[104] Schulzrinne Reply Decl. ¶ 69.

[105] Schulzrinne Reply Decl. ¶ 69.

[106] See, e.g., *'Section 702' Saves Lives*, *supra* note 47; *Section 702 Overview*, *supra* note 49, at 4.

1. That the government's unambiguous concession to the FISC—"*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*"[107]—was false, despite the NSA's representative confirming its accuracy at her deposition.[108]
2. That the PCLOB's statement—that the technical design of the upstream collection program supports the NSA's goal "*to comprehensively acquire communications that are sent to or from its targets*"—was false.[109]
3. That the NSA decided to significantly limit the scope of its collection under the upstream collection program by doing at least one of the following:
    a. Deliberately blacklisting all Wikimedia IP addresses.
    b. Deliberately whitelisting a subset of IP addresses other than Wikimedia's IP addresses.
    c. Deliberately excluding all of the communications protocols that Wikimedia uses, through blacklists or whitelists configured to exclude:
        i. All web activity (i.e., HTTP/S), notwithstanding the government's concession that upstream collection involves the collection of "web activity," *and*

[107] Appendix P (FISC Opinion at 45).

[108] Appendix K (Richards Dep. 160:4-17).

[109] Appendix F (PCLOB Report at 10).

ii. All e-mail activity (i.e., SMTP), notwithstanding the government's concession that, under the upstream collection program, it uses e-mail addresses as selectors.[110]

4. That the NSA has limited the scope of its upstream collection in this way on *every* one of the international Internet links it monitors.

5. That the limitations on collection above are in fact entirely effective at avoiding Wikimedia's communications, even though there are multiple circumstances in which they would not be. For example, Dr. Schulzrinne's claim requires that none of the following could ever occur:

   a. If the NSA whitelisted IP addresses other than Wikimedia's IP addresses:

      i. A user of a whitelisted IP address communicates with Wikimedia, and the communication traverses an international Internet link monitored by the NSA.

   b. If the NSA blacklisted all Wikimedia IP addresses:

      i. One of Wikimedia's communications is enclosed in a multi-communication transaction (MCT) that is not blacklisted, and it traverses an international Internet link monitored by the NSA.

      ii. One of Wikimedia's communications passes through an intermediary that replaces Wikimedia's IP address (such as an e-mail server, VPN, or other communication service)

[110] Appendix F (PCLOB Report at 7).

with an IP address that is not blacklisted, and it traverses an international Internet link monitored by the NSA.

iii. A user of a whitelisted IP address communicates with Wikimedia, and the communication traverses an international Internet link monitored by the NSA.

c. If the NSA excluded the protocols that Wikimedia's communications use:

i. One of Wikimedia's communications is enclosed in a multi-communication transaction (MCT) using a protocol that is not excluded, and it traverses an international Internet link monitored by the NSA.

ii. One of Wikimedia's communications passes through an intermediary using a protocol that is not excluded, and it traverses an international Internet link monitored by the NSA.

6. That the NSA uses an in-line filter architecture to accomplish the upstream collection program in all cases, rather than ever using a copy-then-filter architecture.

155. Given these considerations and my analysis of the NSA's disclosures regarding the upstream collection program, it remains my opinion that it is virtually certain that the NSA has, in the course of the upstream collection program, copied, reassembled and reviewed at least some of Wikimedia's communications.

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

Date: 3/8/2019

Scott Bradner