IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF MARYLAND

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

Exhibit 6

IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF MARYLAND

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

(SECOND) DECLARATION OF DR. HENNING SCHULZRINNE

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

INTRODUCTION AND SUMMARY

1. I am the Julian Clarence Levi Professor of Computer Science at Columbia University in New York, New York. I previously submitted a declaration in this case, dated November 12, 2018. I submit this second declaration at the request of the United States Department of Justice to address the conclusions reached by Mr. Scott Bradner in his December 18, 2018, declaration, including Mr. Bradner's assessment of conclusions reached in my earlier declaration. My background and qualifications in the fields of computer science, electrical engineering, and digital communications technology; the sources of information I considered in arriving at the conclusions stated in this case (apart from those cited herein); and my compensation for my services in this matter, are all stated in my prior declaration.

2. For the reasons I detail herein, it remains my conclusion that the hypothesis advanced in this case by plaintiff Wikimedia Foundation ("Wikimedia"), that the National Security Agency ("NSA"), in the course of conducting Upstream collection, must as a matter of technological necessity be intercepting, copying, and reviewing at least some of Wikimedia's electronic communications that traverse the Internet, is incorrect. Based on what is publicly

known about the NSA's Upstream collection technique, the NSA in theory could be conducting this activity, at least as Wikimedia conceives of it, in a number of technically feasible, readily implemented ways that could avoid NSA interaction with Wikimedia's online communications. Nothing stated in Mr. Bradner's declaration, or in Wikimedia's summary judgment opposition brief ("Wikimedia Brief") alters that conclusion.

3. Mr. Bradner addresses a question quite different from the issue that I was asked to address: the likelihood that the NSA has copied and scanned for targeted selectors at least some of Wikimedia's international text-based Internet communications in the course of conducting Upstream collection. Declaration of Scott Bradner (Plaintiff's Exhibit 1) ("Bradner Decl.") ¶¶ 1(c), 6(e). As both Mr. Bradner and Wikimedia remark, my earlier declaration did not address that issue. Bradner Decl. ¶ 7(a); Wikimedia Brief at 12. Nor do I address it here. That is in part because the Department of Justice has not asked that I opine on that question, which implicates matters that remain classified. I also do not address that question because to arrive at an answer would require that I, like Mr. Bradner, engage in speculation about the NSA's surveillance priorities, practices, and capabilities, and the number and nature of its Upstream surveillance targets, matters about which neither I nor, apparently, Mr. Bradner, has any specialized knowledge or information. I have explained what it is technologically possible for the NSA to do, if it wishes, in conducting Upstream surveillance, and explained how it could avoid interaction (whether by design, or effect) with Wikimedia's communications. But I do not attempt to reach conclusions about what it actually does based on assumptions that lack a basis either in Internet technology or engineering, or pertinent information about the Upstream program.

4. What I do address in this declaration, at the Justice Department's request, are the bases on which Mr. Bradner concludes that it is "virtually certain" that the NSA, in conducting Upstream surveillance, has copied and reviewed at least some of Wikimedia's communications. Bradner Decl. ¶ 6(e). Specifically, I discuss the bases of his conclusions (i) that the NSA, in conducting Upstream surveillance, "most likely" copies, reassembles, and scans for selectors all

communications packets traversing an international Internet link that is monitored by the NSA (if any); (ii) that it is "implausible" that the NSA uses the traffic-mirroring techniques (white- and blacklisting) described in my first declaration; and (iii) that even if the NSA uses one or more of the techniques I described, it is still "virtually certain" that the NSA copies and scans at least some of Wikimedia's communications. As I explain below, Mr. Bradner's conclusions rest principally on assumptions he makes about the NSA's practices and priorities, its resources and capabilities, and its Upstream surveillance targets, and are not based on facts and information concerning Internet technology and engineering. (As was the case with my first declaration, in reaching the conclusions stated herein I have not considered, nor have I been provided with, any classified or other non-public information concerning Upstream surveillance.)

5. It is not the case, as Wikimedia states, that my analysis "ignores key features" of and "critical" disclosures about Upstream surveillance, such as the supposed comprehensiveness of its goals, collection of wholly domestic "about" communications, and the trade-offs of various possible collection system configurations. *See* Wikimedia Brief at 1, 21-22. Rather, these matters were not pertinent to the question of technical feasibility that I was asked to address. And as I explain below, they also provide no support, and certainly none based in Internet technology and engineering, for concluding that the NSA "almost certainly" (Bradner Decl. ¶ 6(a)) copies and scans all communications traversing any circuit it monitors, including Wikimedia's.

6. Finally, at the Justice Department's request, I discuss reasons why, in today's digital communications environment, (i) an organization operating any but the smallest websites would not want to compromise the interests of security, and user privacy, by failing to implement HTTPS encryption "by default," and (ii) an organization that transmits confidential information over the Internet from one business location to another would not want to risk compromising the security and privacy of that information by failing to encrypt it through, for example, the Internet Protocol Security (IPsec) suite.

SUMMARY OF CONCLUSIONS REACHED IN MY FIRST DECLARATION

7. In my prior declaration, I explained that, contrary to Wikimedia's central hypothesis, the NSA could be conducting Upstream surveillance, at least as it is envisioned by Wikimedia, in a number of technically feasible, readily implemented ways that would not require the NSA to electronically copy, electronically scan for selectors (the process Mr. Bradner refers to, less precisely, as "reviewing"), or otherwise interact with *all* communications traversing any given Internet backbone link where Upstream surveillance might theoretically be conducted. I also concluded, more specifically, that if the NSA (hypothetically) were using these techniques to conduct Upstream surveillance, then it could do so without accessing or interacting with Wikimedia's communications. Declaration of Dr. Henning Schulzrinne (Government Exhibit 1) (hereinafter, "First Decl.") ¶¶ 1, 15, 53. These techniques, types of traffic-mirroring known as whitelisting and blacklisting, involve programming a router or switch at a monitored link, using access control lists to selectively mirror (that is, copy) only those communications that are deemed most likely to include communications of interest, without copying or otherwise handling those that are not.

8. For example, if a collecting entity has information indicating that communications of interest are associated with particular IP addresses (or blocks of IP addresses), then at each monitored link an assisting telecommunications service provider could configure its router or switch with a whitelist of the specified IP addresses, which would allow only those communications packets containing source or destination IP addresses on the whitelist to be copied and passed through an interface with collection equipment to be scanned, and (where targeted selectors are detected) retained in the collector's databases. First Decl. ¶¶ 65-66.

9. Conversely, if the collecting entity determines that communications traffic to and from certain IP addresses, or blocks of IP addresses, are of little or no interest for its purposes, then it may request the telecommunications service provider at a given monitored link to configure its router or switch with a blacklist of these IP addresses that would prevent any

packets to or from those addresses from being copied and passed through the interface with the collector's equipment to be scanned. First Decl. ¶¶ 67-68.

10. As I also explained, distinct types of communications, such as e-mail, or web (HTTP and HTTPS) communications, can be whitelisted or blacklisted in the same fashion, based on their assigned port numbers. First Decl. ¶¶ 70-71.

11. Mr. Bradner does not maintain that it is necessary, as a technical matter, to copy all communications traversing a given Internet backbone link that a collecting entity may be monitoring in order to obtain access just to some of them. He repeatedly acknowledges, in fact, that it is possible to copy only a subset of the traffic crossing that link, that is, the communications packets meeting specified criteria such as source or destination IP addresses, or port or protocol numbers. Bradner Decl. ¶¶ 272(b), 280-81, 299, 325, 366.

12. As I also explained, using the traffic-mirroring techniques described in my first declaration, it would be technically feasible for the NSA to conduct Upstream surveillance, at least as it is envisioned in Wikimedia's First Amended Complaint (what I referred to in my first declaration as "Upstream-type" surveillance, *see* First Decl. ¶¶ 15, 77, 88), without copying, scanning, or otherwise interacting with any of the three categories of communications that Wikimedia believes are subject to NSA Upstream collection processes. First Decl. ¶ 77. The first category includes HTTP and (principally) HTTPS requests from individual Internet users to the servers housing Wikimedia's websites, and the responses thereto. First Decl. ¶ 78. These communications could be blacklisted by configuring the router or switch at any monitored link to prevent any communications with port numbers 80 and 443, respectively, from being copied and forwarded to NSA scanning equipment. First Decl. ¶ 79. Alternatively, Wikimedia's HTTP and HTTPS communications could be excluded from those made available to the NSA at a monitored link through (i) blacklisting communications to or from Wikimedia IP addresses, or (ii) using whitelists of IP addresses (that do not include Wikimedia's IP addresses) to exclude Wikimedia HTTP and HTTPS communications (and other communications of no intelligence interest) except

those exchanged with users (if any) who had been assigned targeted (whitelisted) IP addresses. First Decl. ¶¶ 80-81.[1]

13. Similarly, NSA access to Wikimedia's second category of communications, encrypted log communications, could be blocked by blacklisting communications containing the protocol number (50) of the IPSec protocol used by Wikimedia to encrypt its log communications, or, as in the case of Wikimedia's Category 1 communications, whitelisting or blacklisting by IP address. First Decl. ¶¶ 83-84. The third category, various types of online communications engaged in by Wikimedia's staff, could also be blocked, in the same manner as Wikimedia HTTP and HTTPS communications, through white- or blacklisting by IP address. First Decl. ¶ 87.

MR. BRADNER'S ASSERTION THAT THE NSA "MOST LIKELY"CONDUCTS UPSTREAM SURVEILLANCE USING HIS COPY-ALL-THEN-SCAN APPROACH IS WITHOUT A BASIS IN INTERNET TECHNOLOGY AND ENGINEERING

14. Although Mr. Bradner acknowledges the technical feasibility of the white- and blacklisting techniques I describe in my earlier declaration, he states that in his view it is "most likely" that the NSA does not use these filtering techniques, and instead copies and (after reassembling the packets) scans for selectors all communications, including Wikimedia's, that traverse any international Internet backbone link (if any) where the NSA conducts Upstream surveillance. Bradner Decl. ¶¶ 282-89, 366-67.

15. For the most part, Mr. Bradner's conclusions regarding the "likely" manner in which the NSA actually conducts Upstream collection are based on assumptions about the NSA's policies and practices in conducting covert surveillance, its operational priorities, and resources, and the number of its Upstream surveillance targets. As I stated above, I have not been asked to offer views on the "likely" manner in which the NSA actually conducts Upstream collection, because of its classified nature, and I do not venture an opinion on that subject for the very reason that an informed opinion would require knowledge and information about the classified

[1] I further explained that whitelisting and blacklisting by IP address in this fashion would also block NSA access to the SMTP (e-mail) communications, of unspecified number and geographic distribution, that Wikimedia also includes in Category 1 of its communications. First Decl. ¶ 82.

operational details of the Upstream program that I (and by all appearances, Mr. Bradner) lack. To a lesser extent, Mr. Bradner attempts to rest his conclusions on what he assumes would be the preferred equipment configuration of an assisting telecommunications service provider, but ultimately this assumption rests on a misreading of my earlier declaration, and further speculation regarding the number of the NSA's Upstream targets.

16. Mr. Bradner considers it likely that an assisting provider at any monitored link furnishes the NSA with copies of all communications packets traversing that link, after which (in Mr. Bradner's opinion) the NSA reassembles the packets and then reviews (electronically scans) the reconstructed communications for targeted selectors in order to identify those that will be retained in NSA databases. Bradner Decl. ¶¶ 273-79, 282-89. Mr. Bradner refers to this process as a "copy-then-filter" configuration." Bradner Decl. ¶ 272(a). For clarity of distinction, I will refer to this approach herein as "copy-all-then-scan." Mr. Bradner considers this approach more likely than use of the traffic-mirroring techniques described in my first declaration, which he refers to as "in-line filtering," Bradner Decl. ¶¶ 272(b), 280-89, and which I will refer to herein (again for clarity of distinction) as "filter-then-copy-and-scan."

17. Mr. Bradner gives four reasons for considering his copy-all-then-scan approach more likely than the filter-then-copy-and-scan techniques described in my first declaration. The first two concern the supposed operational preferences of the NSA, and the latter two the supposed preferences of an assisting provider. The reasons Mr. Bradner gives are:

> a. that whitelisting or blacklisting (by IP address or port or protocol number) before copying communications and scanning them for targeted selectors would require the NSA to share sensitive information about its targets and/or filtering criteria with the assisting provider's personnel. Bradner Decl. ¶¶ 283, 285-87.
>
> b. that whitelisting or blacklisting to reduce the volume of communications that must be reassembled and scanned for selectors would be of little offsetting benefit to the NSA given the real-time processing capacity of modern packet-inspection devices. Bradner Decl. ¶ 288.
>
> c. that his suggested copy-all-then-scan configuration would not require the placement of "an NSA-operated device into the heart of [the provider's] network," which would risk an adverse impact on the provider's network "in the event of an equipment failure or misconfiguration." Bradner Decl. ¶ 284.

d. that configuring a router or switch to filter the communications made available to the NSA could create a risk of "overloading" the router and impairing the provider's ability to support its customers' traffic. Bradner Decl. ¶¶ 288, 366(c).

I address these purported justifications in turn.

18. Sharing Sensitive Information: The extent to which the NSA is willing (or finds it necessary) to share classified information with an assisting provider in order to conduct Upstream surveillance (or any other kind of collection activity) is a matter about which neither I nor (so far as his declaration reveals) Mr. Bradner has any specialized knowledge or information. I do not consider uninformed assumptions about the NSA's willingness to share such information with an assisting provider to be a basis on which to reach conclusions, from the perspective of Internet technology and engineering, about the manner in which the NSA conducts Upstream surveillance. I observe, however, that according to the PCLOB Section 702 Report cited by Mr. Bradner and Wikimedia, Privacy and Civil Liberties Oversight Board, *Report on the Surveillance Program Operated Pursuant to Section 702 of the Foreign Intelligence Surveillance Act* (July 2, 2014), https://www.pclob.gov/library/702-Report.pdf, the NSA already shares sensitive information about its surveillance targets with assisting provider(s), specifically, the selectors (such as their e-mail addresses and telephone numbers) used to identify their communications for acquisition. PCLOB Section 702 Report at 36. This would appear to call Mr. Bradner's premise into question.

19. Operational Benefit of Filtering Traffic Before Scanning: Similarly, whether the NSA places greater importance on the potential intelligence value of scanning every communication that crosses a given Internet backbone link, or the operational efficiencies and cost-savings that would flow from first filtering out communications of low interest, is a matter about which neither I nor, evidently, Mr. Bradner, has knowledge. That said, the practical benefits to be gained from first filtering out low-interest communications cannot be dismissed on the grounds suggested by Mr. Bradner. Mr. Bradner simply asserts, without supporting data or explanation, that there is little to be gained by reducing the volume of traffic that must be copied and scanned for selectors because (in his estimation) the most powerful, commercially

available packet-inspection devices are capable of scanning all traffic in real time as it traverses a given monitored link. *See* Bradner Decl. ¶ 288 ("Modern deep packet inspection devices, individually or operating in parallel, can process or review Internet communications at the same rate that those communications traverse high-bandwidth Internet links."). (He also implicitly assumes that these high-end commercial packet-inspection devices, or devices of similar capacity, are those actually employed by the NSA.).

20. Readily available data concerning traffic volume at the links that Wikimedia claims are monitored by the NSA, and the processing capacity of the packet-inspection devices that Mr. Bradner refers to, contradict his assumptions. The data transfer rates (i.e., the traffic flow) at international Internet links of the kind that Wikimedia presumes are monitored by the NSA are likely to surpass the processing capacity of even the best-resourced entities. For example, the AEConnect link, a transatlantic cable running between Shirley, New York and Killala, Ireland, put into service in 2016, has an aggregate capacity of approximately 40 terabits (over 40 trillion bits) of information per second[2]. The newer MAREA link, a suboceanic cable connecting Virginia Beach, Virginia with Bilbao, Spain, has a capacity of 160 terabits (over 160 trillion bits) of information per second.[3] According to the telecommunications market research firm TeleGeography, inter-regional Internet capacity has increased to 98 terabits per second, while total international capacity (a large fraction of which is known to originate or terminate in the United States) reached 295 terabits per second in 2017[4]. In contrast, the largest commercial deep-packet-inspection (DPI) devices, such as the NIKSUN Supreme Eagle, typically have Ethernet interfaces with speeds of no more than 100 gigabits (100 billion bits) per second,[5] that is, 1/10 terabit per second, as those are the fastest standardized Ethernet interfaces that are commercially available.

[2] https://www.submarinenetworks.com/systems/trans-atlantic/aeconnect

[3] https://en.wikipedia.org/wiki/MAREA

[4] https://blog.telegeography.com/295-tbps-internet-traffic-and-capacity-in-2017

[5] https://www.niksun.com/c/1/ds/NIKSUN_datasheet_Supreme_eagle.pdf

21. Not all transoceanic fiberoptic cables, especially those more recently put into service, operate at full capacity. But even at half capacity (20 terabits per second for AEConnect, 80 terabits per second for MAREA), to copy and scan all traffic crossing the AEConnect link, as Mr. Bradner suggests, would require 200 packet-inspection devices, and at the MAREA link 800 such devices, with racks upon racks full of monitoring gear. In addition, a collection system such as Mr. Bradner envisions would require installation of an opto-electronic device, such as a router, to convert the single high-speed stream of traffic flowing over the link into hundreds of lower-speed streams feeding into the individual packet-inspection devices. Acquiring, deploying, operating, and maintaining one or more collection systems of this kind would be enormously costly, and present formidable technical and logistical challenges. Thus, there are potentially many reasons to have the assisting carrier's router perform the preliminary filtering, removing most of the traffic that is likely not of interest before communications are copied, reassembled, and scanned. (For example, removing encrypted web traffic alone would likely reduce the volume to be copied and scanned by about half.)

22. Mr. Bradner also makes the related observation that "[i]f filtering traffic for performance reasons were desirable, the NSA would get much more result from filtering YouTube than from filtering Wikimedia." Bradner Decl. ¶ 366(c). That observation ignores the fact that the volume of streaming video traffic at international links is far lower than at domestic links, since most video streaming services, such as Netflix, do not transmit programming across international boundaries. (Most viewers access streaming video services, such as Netflix, YouTube, or Hulu, from content distribution networks located within their home countries, or at least their own continents.) Therefore, the benefit of filtering out only video traffic crossing an international Internet link (where Wikimedia supposes that Upstream surveillance is conducted) would likely be modest. Moreover, Mr. Bradner's observation is a non-sequitur; the two options, filtering video traffic, and filtering Wikimedia, are not mutually exclusive. If the NSA wished to avoid scanning large volumes of communications traffic of (hypothetically) no intelligence

interest, then there is no technological reason why it could not filter both video traffic *and* traffic from high-volume websites such as Amazon.com and Wikipedia.

23. Placement of NSA Equipment in a Provider's Network: I turn now to the reasons given by Mr. Bradner for his view that a provider would prefer his copy-all-then-scan configuration to a filter-then-copy-and-scan approach. The central premise of this conclusion is that the filter-then-copy-and-scan approach I describe in my first declaration involves placement of "an NSA-operated device into the heart of [a provider's] network." Bradner Decl. ¶ 284. That assertion is simply mistaken. I did not propose in my earlier declaration a filtering configuration involving the use of "NSA-operated devices." The whitelisting and blacklisting techniques I described all involve the use of existing provider equipment, that is, the provider's own network switch or router, programmed with access control lists according to NSA criteria, in order to perform the very kind of traffic-mirroring function for which routers and switches are designed and utilized in the ordinary course of a provider's business. First Decl. ¶¶ 65-71.

24. Moreover, the techniques I described would pose little risk to the operational integrity of the provider's network, i.e., the unimpeded flow of communications traffic to and from the provider's customers. Adding monitoring capability, with the constraints noted, to an existing carrier router does not affect carrier operations, requires minimal physical additions to carrier facilities (such as running an extra fiber patch cable) and operationally speaking is imperceptible to the carrier. It only requires periodic configuration (programming) changes in the white- and blacklists that are stored electronically in the router's memory, which can be accomplished (either by NSA or provider personnel) through remote rather than onsite access, not unlike the way a document stored on a server can be updated and edited by a user who connects to the server from another location.

25. Risk of "Overloading" the Provider's Router: When Mr. Bradner states that whitelisting or blacklisting "could affect the performance of the [provider's] router and create a risk of overloading the router," Bradner Decl. ¶ 288, he is again making assumptions, this time about the extent to which filtering communications made available to the NSA would place

demands on a router's processing capacity. That would depend in large part on the number of IP addresses (or blocks of addresses) that the NSA designates for white- or blacklisting at a given monitored link, which in turn would be related, to a significant degree, to the number of the NSA's Upstream targets (and to the number of targets whose communications the NSA is seeking to capture at that particular link). These are also matters about which Mr. Bradner has no knowledge or information, at least so far as his declaration reveals.

26. There are, in addition, a number of other practical considerations, overlooked by Mr. Bradner, that weigh against choosing Mr. Bradner's copy-all-then-scan approach over a filter-then-copy-and-scan approach. There are two possible equipment configurations for implementing Mr. Bradner's suggested approach, (a) the use of an optical splitter to create a copy of the entire communications stream flowing across a monitored circuit, *see* First Decl. ¶ 55, or, alternatively, (b) configuring the provider's router to mirror all incoming or outgoing packets to the collection infrastructure. Both pose risks of adverse impacts on the network to which a provider might object.

27. When an optical splitter is connected to an optical fiber, all communications traffic carried on that fiber is directed into the splitter, where the stream is duplicated by dividing the optical power of the stream between two (or more) outputs, allowing the "original" communications stream to continue to its intended destination, albeit at reduced optical power, while the duplicate stream(s) may be diverted for other purposes, including scanning for communications of interest. Thus, while passive optical splitters are relatively simple components, adding a splitter to facilitate Upstream collection would introduce another potential failure point to a provider's network, and at best introduce a degree of optical power loss. Generally, to ensure high reliability and easy maintenance, optical network architectures try to minimize both the number of connectors and optical loss.

28. As an alternative to optical splitting, a router could theoretically be configured to mirror all incoming or outgoing packets to the monitoring infrastructure. However, no commercial router I am aware of is designed to mirror all incoming or outgoing traffic at once.

Rather, carrier-scale routers may only be able to mirror traffic to one or two interfaces by design. To mirror all traffic (if at all possible) would require adding interfaces capable of supporting the total input or output capacity of the router, potentially exceeding its design capacity and doubling its cost. (Interface cards constitute the largest single cost component of a carrier-grade router. Each router model has a maximum interface capacity that constrains the number and speed of interfaces.)

29. I do not mean by these observations to suggest how the NSA or the assisting provider at a given link theoretically monitored by the NSA would choose to configure an Upstream collection system. I am saying, however, that the reasons given by Mr. Bradner for suggesting that the NSA and the provider would prefer his copy-all-then-scan configuration over a filter-then-copy-and-scan approach (i) make assumptions about the NSA's surveillance practices, priorities, and resources, (ii) are based on a misunderstanding of the configuration I proposed in my earlier declaration, (iii) rely on assumptions, rather than information, concerning the number of the NSA's Upstream targets, and (iv) overlook important considerations that a provider concerned about network integrity would take into account. I do not consider them a reliable technological basis on which to conclude that it is "most likely" the NSA is using a copy-all-then-scan an approach, as asserted by Mr. Bradner.

MR. BRADNER'S VIEW THAT USE OF A FILTER-THEN-COPY-AND-SCAN APPROACH IS "IMPLAUSIBLE" ALSO LACKS A BASIS IN INTERNET TECHNOLOGY AND ENGINEERING

30. Conversely, Mr. Bradner gives a number of reasons why he concludes that the filter-then-copy-and-scan approach described in my first declaration is "implausible." Bradner Decl. ¶¶ 366-67. A number of these reasons are the same as those he gives in support of his conclusion that the NSA "most likely" uses his copy-all-then-scan approach. *See* Bradner Decl. ¶ 366(a) (providing sensitive information to the assisting provider); ¶ 366(c) (no need to reduce the processing load on packet-inspection devices). I have already addressed the conjectural nature of these points and why, as a technological matter (that is, from the perspective of Internet technology and engineering) they do not support Mr. Bradner's conclusions. Below I

address the remaining bases of Mr. Bradner's "plausibility" conclusion, which may be divided into three categories: (a) those concerning port and protocol blocking (Bradner Decl. ¶ 366(b), (e)-(h)); (b) those concerning whitelisting (Bradner Decl. ¶ 366(d)); and (c) those concerning blacklisting Wikimedia IP addresses (Bradner Decl. ¶ 367(a)).

Blocking Port or Protocol Numbers (HTTP and HTTPS communications)

31. Mr. Bradner first considers it implausible that the NSA would blacklist particular types of communications by port or protocol number, in particular HTTPS communications (encrypted web communications), for a variety of reasons. Bradner Decl. ¶ 366(b), (e)-(h). As I now explain, each ground given by Mr. Bradner for this conclusion is based on assumptions about the NSA's surveillance priorities, or is unexplained altogether, and provides no technological basis for concluding that the NSA would not (much less could not) utilize this filtering technique.

32. Principally, Mr. Bradner remarks that blocking particular types of communications by port or protocol number would leave "blind spot[s] in the NSA's Upstream surveillance that "[s]ophisticated targets" could "easily probe" to discover and exploit to avoid collection of their communications. Bradner Decl. ¶ 366(b), (e). Mr. Bradner does not explain what targets could "probe," or how, to discover these so-called blind spots. The traffic-mirroring techniques I describe in my earlier declaration are completely invisible to users' end systems (that is, their communicating devices) and to other equipment on the network not directly engaged in handling the monitored traffic. For example, traffic mirroring does not increase the delay or reduce the data flows being mirrored, or change the content or headers of the packets being transmitted in any way. Even if the undefined "probing" Mr. Bradner alludes to were possible, he does not explain what level of technical sophistication would be required, or on what basis he assumes that the NSA's targets possess that level of sophistication.

33. Because port numbers are, in Mr. Bradner's words, "only advisory," Bradner Decl. ¶¶ 109, 366(e), he suggests that if potential NSA targets somehow found out that the NSA was (hypothetically) blacklisting ports 80 and/or 443 (in order to block HTTP and/or HTTPS communications from its collection devices), then these potential targets could assign port 80 or

443 to all their communications (whether in fact they are HTTP or HTTPS communications or not) and thereby avoid detection. Bradner Decl. ¶ 366(b), (e). While port numbers cannot be dismissed as merely "advisory,"[6] it is technically feasible, with certain pre-arrangements, for two or more users to communicate in the manner that Mr. Bradner describes. At bottom, however, whether the creation of "blind spots" is a matter of such genuine intelligence concern as to motivate the NSA to examine all HTTP and HTTPS communications (even if it were not otherwise persuaded of the value in doing so), depends on facts and information concerning its mission priorities and resources known only to the NSA (or at least not presented in Mr. Bradner's declaration).[7]

34. Apart from the "blind spot" issue, Mr. Bradner also remarks that blacklisting HTTP and HTTPS communications (ports 80 and 443) "would leave a very large hole in the NSA's collection ability," including web-based e-mail, webchat, and web-based editors, and that there are many "obvious" reasons, in his view, for the NSA to acquire HTTPS communications. Bradner Decl. ¶ 366(f), (g); *see also* Bradner Decl. ¶¶ 326, 359. In so contending, Mr. Bradner is apparently making assumptions about the value that the NSA places on particular types of communications rather than offering a technological reason why the NSA could not or would not block access to such communications.

[6] Port numbers are assigned to specific applications, such as e-mail or web browsing, by the global Internet Assigned Numbers Authority (IANA), *see* First Decl. ¶¶ 16, 36, and represent established international conventions—default settings, in a manner of speaking—to facilitate automated, "user-friendly" communication between devices connected to the Internet. While use of an assigned port to run an application is neither legally nor even technically mandated, using non-standard ports, as I have explained, requires that users on both ends of an exchange must agree in advance to communicate in this atypical fashion, so that appropriate adjustments can be made to their communications before they are transmitted, to ensure that they are routed on receipt to the agreed-upon, non-standard port.

[7] Mr. Bradner observes that if the NSA blacklisted only HTTPS communications, but not HTTP communications, then it could still obtain access to unblocked HTTP communications to and from Wikimedia servers. Bradner Decl. ¶ 366(h). While true in theory, I have not posited a scenario in which the NSA blocks only HTTPS, but not HTTP communications. I have observed only that it is technically feasible to block both, First Decl. ¶¶ 80-81, a point with which Mr. Bradner does not disagree.

35. Moreover, I observe that if the NSA were interested in communications to and from particular websites, such as webmail sites, or chatroom sites, then it would not be necessary for the NSA to obtain access to all HTTP or HTTPS communications traversing a monitored link in order to do so. It is technically feasible, using a combination of blacklisting and whitelisting, to provide the NSA with access only to communications with websites of particular interest. Specifically, at a monitored link the provider's router or switch could be configured with a blacklist that would block NSA access to all communications with port numbers 80 or 443 (i.e., all HTTP and HTTPS communications), except those HTTP and HTTPS communications to or from the IP addresses included on a whitelist containing the addresses of the sites of interest to the NSA (including, hypothetically, specific webmail and chatroom sites). In this fashion, blacklisting HTTP and HTTPS communications (including Wikimedia's) would not necessarily, at least as a technological matter, carve out so large a "hole" in the NSA's Upstream collection as Mr. Bradner assumes.

36. The additional reasons given by Mr. Bradner for concluding that the NSA is "likely" acquiring HTTPS communications include (a) that the NSA is authorized to collect encrypted communications under the "minimization procedures" that govern its surveillance activities under Section 702 of the Foreign Intelligence Surveillance Act ("FISA"), Bradner Decl. ¶¶ 325, 366(g); (b) that the NSA has acknowledged collecting "web activity" under Section 702 of FISA, Bradner Decl. ¶¶ 314-15, 366(f) & n.126; (c) that the NSA "may, currently or in the future, be able to decrypt important encrypted messages," Bradner Decl. ¶ 326(a), (b); and (d) that the NSA could learn "useful information," such as the IP addresses of potential targets and the websites they visit, from the unencrypted addressing information of encrypted HTTPS communications, Bradner Decl. ¶¶ 326(c), 366(g). None of these reasons is a technological basis on which to assess the likelihood that the NSA does or does not acquire HTTPS communications, and, even on their own terms, they do not support Mr. Bradner's conclusion.

a. Authority under Section 702 to collect encrypted communications: There are a variety of encrypted Internet applications and communications, using a variety of

encryption techniques, apart from HTTPS, and it is my understanding that the NSA conducts at least two forms of collection under authority of FISA Section 702 – PRISM and Upstream. Therefore, to say that the NSA is authorized generally to collect encrypted communications under Section 702, Bradner Decl. ¶¶ 325, 366(g), is not to say specifically that it acquires HTTPS communications using Upstream collection.

b. Collection of "web activity": The reference to "web activity" in the NSA court filing cited by Mr. Bradner, Bradner Decl. ¶¶ 314-15, 366(f), App'x C at 30, appears in a passage comparing the total number of so-called multi-communication transactions (MCTs) acquired through Upstream to "the *total* take of Section 702 upstream collection of web activity" (emphasis mine). This juxtaposition, together with the reference to "total" take, suggests that the term "web" activity may have been intended here to refer to Internet activity as a whole, in light of the fact that MCTs are not necessarily limited to HTTP and HTTPS communications but could include e-mail (SMTP) communications. Moreover, it is common in colloquial usage to use the term "web" when referring to the Internet at large.[8] Even assuming that "web activity" as used in the cited NSA filing refers more precisely to HTTP and HTTPS communications, there are various forms of web activity, such as webmail, and chatrooms, that do not include communications with what we ordinarily think of as websites such as Wikimedia's, and that could be obtained using the combined blacklisting/whitelisting technique I discuss in paragraph 35, above.

c. Possibility of decryption: While the NSA "may" be able, now or in the future, to decipher encrypted communications, Bradner Decl. ¶ 326(a), (b), equally so, it may not, and without information about the NSA's true decryption capabilities, I do not see how the abstract possibility makes it "likely" as a technological matter that the NSA, in fact, collects encrypted communications at all, much less HTTPS communications specifically.

[8] *See* https://en.wikipedia.org/wiki/Internet#World_Wide_Web (observing that "[m]any people use, erroneously, the terms Internet and World Wide Web, or just the Web, interchangeably . . .").

d. Potential intelligence value of addressing information: Mr. Bradner's suggestion that the NSA might find intelligence value in the addressing information of HTTPS communications that would merit their collection even though their contents are encrypted, Bradner Decl. ¶¶ 326(c), 366(g), is simply another assumption on his part about the NSA's surveillance priorities that provides no technological support for his conclusion.

37. However, even taking as given Mr. Bradner's assumption that the NSA, for the reasons above, may be motivated to collect certain HTTPS communications, there are technical means by which it could obtain access to HTTPS communications of interest without copying and scanning all HTTPS communications at a monitored link, including Wikimedia's. As I explained in paragraph 35, above, it is technically feasible, to blacklist ports 80 and 443 (i.e., HTTP and HTTPS communications), while simultaneously whitelisting the IP addresses of websites, webmail services, and/or chatrooms of interest. In this fashion the NSA could obtain access to communications to and from websites of interest while excluding all others, including, hypothetically, Wikimedia's. A configuration of this kind would be entirely consistent, moreover, with the references to acquiring encrypted communications in the NSA's minimization procedures, and to collection of "web activity" in the NSA court filing cited by Mr. Bradner.

38. In short, Mr. Bradner gives no technological reason, as opposed to conjecture about the NSA's practices, priorities, and capabilities, for dismissing port blocking as "implausible."

Blacklisting Wikimedia's IP Addresses

39. Next, Mr. Bradner finds it "basically inconceivable" that the NSA would blacklist Wikimedia's IP addresses, for two reasons. Bradner Decl. ¶ 367(a). Neither is based in Internet technology and engineering.

40. Mr. Bradner states that it is "totally unbelievable" that the NSA would have undertaken the "incredibly resource-intensive task" of sifting through millions of websites to decide which to monitor and which not. Bradner Decl. ¶ 367(a). This is another non-sequitur, as

I did not suggest anything of the kind. What I discussed in my first declaration was the practical possibility of blacklisting certain high-volume but perhaps low-interest websites, such as, hypothetically, Amazon.com, Wikimedia's websites, and perhaps others, First Decl. ¶ 81, in order to reduce unwanted volumes of communications and enhance the efficiency of the collection process. This would be a trivial task.

41. Identifying and removing high-volume, commercial websites and video services is made much easier since page views follow a Zipf distribution, that is, a website's volume of page views declines rapidly as its popularity ranking decreases. For example, according to the Parse.ly blog post,[9] the top roughly 1,000 websites account for 97% of global page views. Thus, while the identity of popular websites shifts over time, these can be tracked relatively easily by widely-known services like Alexa.com, https://www.alexa.com/topsites.[10] According to Alexa, as of early 2019 Wikipedia.org is ranked as the fifth most popular website globally (sixth in the United States), so even a manual process would likely include it on any filtering blacklist (assuming the NSA deemed it of low interest). If desired, the list of these popular sites could be obtained periodically and mechanically, converted to IP addresses by domain name lookups programmatically, and then be used to modify the filter list used in routers.[11] When Mr. Bradner disparages the idea of reviewing millions of websites to decide one by one which sites to monitor and which to ignore, he is taking issue with a proposal that I have not made.

42. Second, Mr. Bradner states that he finds blacklisting Wikimedia IP addresses to be implausible because doing so "would deliberately limit the possible collection of information on the use of Wikimedia resources by NSA targets, a potentially valuable source of information

[9] https://blog.parse.ly/post/10/zipfs-law-of-the-internet-explaining-online-behavior/

[10] Alexa.com tracks the rankings of the popular websites; while also operated by Amazon, it is not related to Amazon's smart-speaker service with the same name.

[11] As Mr. Bradner notes, the NSA, at least in some circumstances, uses IP address filtering to "eliminate potential domestic transactions" from those scanned. PCLOB Section 702 Report at 37; Bradner Decl. ¶ 290. As he also notes, the number of domestic U.S. IP addresses exceeds 60,000. Bradner Decl. ¶ 229. If the NSA can maintain a list of over 60,000 address blocks, "no easy task," Bradner Decl. ¶ 298, then it can also add a few dozen or a hundred IP addresses of organizations that contribute high volumes of traffic, but are unlikely to be of intelligence interest, such as (hypothetically) Wikimedia.

about the online research and reading of its targets." Bradner Decl. ¶ 367(a). Once again, Mr. Bradner is engaging in speculation, without an evident basis in knowledge or information, about the online reading habits of the NSA's targets and the intelligence value that the NSA would ascribe to knowing whether they read Wikimedia websites. That is not a technological basis on which to judge the plausibility of blacklisting Wikimedia websites.

Whitelisting IP Addresses of Interest

43. As I have explained, First Decl. ¶ 65, whitelisting is a filtering technique whereby an assisting telecommunications carrier at an Internet backbone link hypothetically monitored by the NSA could provide the NSA only with copies of communications packets whose source or destination IP addresses appear on a list of particular IP addresses, or blocks of IP addresses, that the NSA has determined are associated with communications (or targets) of interest. Using this technique, the NSA would not obtain access to any of Wikimedia's communications unless users communicating with Wikimedia's websites, or perhaps with its staff, were assigned IP addresses on the targeted whitelist. First Decl. ¶¶ 81, 84, 87.

44. Mr. Bradner suggests no reason why whitelisting, if employed, would not block NSA access to Wikimedia's communications (unless, hypothetically, Wikimedia were receiving or responding to communications from a whitelisted IP address). Rather, he gives a single reason why he finds it implausible that the NSA would employ whitelists in the course of Upstream surveillance. Bradner Decl. ¶ 366(d). Separately, Wikimedia, in its legal brief, suggests that whitelisting would be inconsistent with certain features of Upstream collection. Wikimedia Brief at 22. I address both points in turn.

45. Mr. Bradner opines that whitelists would be "useless" for Upstream collection purposes because it is not "remotely possible" for the NSA to know in advance the IP addresses of all its targets. Bradner Decl. ¶ 366(d). He premises this conclusion, however, on various assumptions about the number, nature, and mobility of the NSA's Upstream surveillance targets, for which he offers no supporting information.

46. First, Mr. Bradner appears to assume that the NSA has over 129,000 Upstream targets, based on a Government report disclosing that in 2017 the NSA had over 129,000 Section 702 targets. As discussed above, however, and as Mr. Bradner also observes, the NSA operates two collection programs under Section 702 of FISA, Upstream and PRISM. *See* Bradner Decl. ¶ 334 & n.108. So far as I am aware, the Government has not disclosed any information concerning the number of Section 702 targets that are targets, specifically, of Upstream collection, as opposed to those that may be targets exclusively of PRISM collection. Mr. Bradner presents no basis on which to assume that the NSA's Upstream targets are too numerous to make whitelisting technically practical.

47. Second, Mr. Bradner simply assumes both (i) that the NSA's Upstream targets include individuals (as opposed to other more stationary types of entities), and (ii) that their movements result in changes to the IP addresses associated with their communications. Bradner Decl. ¶¶ 334, 366(d). Even if the NSA's Upstream targets include individuals who move from place to place, IP addresses, as Mr. Bradner acknowledges, are often associated with specific geographic areas to a "reasonable degree of certainty." Bradner Decl. ¶ 296; *see also* First Decl. ¶¶ 32-34. Therefore, even if an individual NSA target moved to a certain degree from place to place within a given geographic area, in principle the NSA could reliably obtain access to that target's communications by targeting not just a single IP address, but a set of IP addresses, as I posited in my first declaration, First Decl. ¶ 67, associated with geographic areas where the target is believed to be located.

48. In the final analysis, the NSA's capabilities to ascertain and track the IP addresses of communications with its Upstream targets are unknown, and a blanket assumption that it lacks the ability to do so is not a technological basis on which to dismiss whitelisting as "useless."

49. Wikimedia separately argues in its legal brief that the idea of whitelisting "ignores" the NSA's collection of "about" communications that are neither to nor from its targets, and the fact that NSA selectors (such as e-mail addresses) do not appear in the packet headers where

network-layer addressing information, such as IP addresses, is located. Wikimedia Brief at 21-22. Notably, this is not a point echoed by Mr. Bradner, and it is not well taken.

50. Wikimedia appears to be conflating two separate steps in the filter-then-copy-and-scan collection process I have described. First is the filtering of communications within the provider's router or switch as they traverse a monitored link. In the whitelisting context, this step involves a comparison of packets' source and destination IP addresses (contained in the packet header) to a whitelist of specified IP addresses, so that only packets containing source or destination IP addresses on the whitelist are copied and made available for scanning. The second step occurs after the whitelisted packets are copied and passed through the router interface to be scanned by the collection system. This is the point at which the packets' "payload," including, for example, source and destination e-mail addresses as well as their message content, is scanned for targeted selectors, to identify those communications that will actually be retained.

51. This two-step process is entirely compatible with the acquisition of "about" communications as described in public sources, and with scanning for selectors other than IP addresses. If Wikimedia means to suggest otherwise, then it is mistaken. If the NSA, through whitelisting, were to obtain access to communications to and from IP addresses of interest, it could then scan them for the presence of targeted selectors of any (authorized) kind, be they e-mail addresses, telephone numbers, or (hypothetically) other communications identifiers. And if the NSA obtained access to communications to or from a specified (whitelisted) set of IP addresses it has associated with a target (or a target's geographic location), and scanned them for the target's e-mail address, then it could acquire not only communications to or from the target's e-mail address, but also "about" communications between parties other than the target, also originating from or destined for one of the whitelisted IP addresses, that contain the target's e-mail address in their contents.

52. In sum, neither Mr. Bradner, nor Wikimedia, has identified any grounds based in Internet technology and engineering that would render whitelisting "useless," or otherwise impractical or "implausible" for purposes of NSA Upstream collection. They have provided no

reason to retreat from my conclusion that at any Internet backbone link that the NSA might hypothetically be monitoring, the NSA, through whitelisting, could block access to any Wikimedia communications, except those (if any) in which users communicating with Wikimedia, had been assigned a targeted IP address. First Decl. ¶¶ 81, 84, 87.

MISCELLANEOUS POINTS RAISED BY MR. BRADNER AND WIKIMEDIA

53. In the foregoing sections I have addressed (i) the grounds on which Mr. Bradner concludes it is "more likely" that the NSA employs his copy-all-then-scan approach to Upstream collection, rather than a filter-then-copy-and-scan approach, and (ii) the grounds on which Mr. Bradner (and, separately, Wikimedia) argue that a filter-then-copy-and-scan approach using whitelisting and/or blacklisting techniques is "implausible." In this section I turn to several more general observations made by Mr. Bradner, and Wikimedia.

54. Asymmetric routing of Internet communications: Mr. Bradner remarks that my analysis "does not discuss the asymmetric routing of communications on the Internet." Bradner Decl. ¶ 359. To the contrary, as Mr. Bradner and I both have observed, communications exchanged between two users on the Internet (such as e-mail), or communications between a user and a website, do not necessarily follow the same path back and forth, even in real time. Bradner Decl. ¶¶ 104(b), 199, 309; First Decl. ¶ 89. I believe that Mr. Bradner overstates the impact of asymmetric routing between two end points at international gateways. (The number of international gateways between any two countries is relatively small.) But the point is that asymmetric routing has nothing to do with the feasibility or utility of whitelisting and/or blacklisting as (hypothetical) Upstream collection techniques that could block NSA access to Wikimedia communications.

55. As Mr. Bradner himself observes, the upshot of asymmetric routing, when it occurs, is that the response to a communication may cross a different international link than the one crossed by the original communication, and that the NSA, if interested in acquiring both the original communication and the response, would have to monitor both links. Bradner Decl. ¶ 309. This has no impact, however, on the scope of the communications the NSA must monitor

at each link, regardless of the number of links actually monitored. For example, if the source IP address of the original communication were included on a whitelist used at the first link, then, if the response returned over a different link, it could still be captured at that second link, using the same whitelist, since the source IP address of the original communication would now appear as the destination IP address of the response. Thus, it would not ordinarily be necessary for the NSA to copy and scan all communications, at either link, in order to capture both the original communication and the response.

56. Acquisition of wholly domestic "about" communications: Mr. Bradner suggests that the NSA must not be using IP address filtering to eliminate wholly domestic communications before copying and scanning, at least at some monitored links. Bradner Decl. ¶ 293. He bases this conclusion on a statement, in an October 2011 opinion by the Foreign Intelligence Surveillance Court ("FISC"), that the NSA has acknowledged that it "will acquire a wholly domestic 'about' communication" if it is "routed through an international internet link being monitored by the NSA." Bradner Decl. ¶¶ 292-94 & n.83, App'x P. Whether or not that is so, it would be technologically inaccurate to conclude on this basis (as Wikimedia appears to do, Wikimedia Brief at 21-22), that the acquisition of some wholly domestic communications, even at a so-called "international Internet link," is inconsistent with the use of the whitelisting and blacklisting techniques I have described.

57. As Mr. Bradner observes, the "routing of wholly domestic communications over international circuits does occasionally happen." Bradner Decl. ¶ 292 & n.82. There are several scenarios in which this could occur. For example, suppose that a person located in the United States, who uses a foreign-based virtual private network (VPN) service, sends an e-mail to another person, also located in the United States. When a user communicates via a VPN, all of the user's communications are encrypted and first routed through the VPN server before being directed to their ultimate destination. As a result, on that first leg each communications packet

is assigned the VPN server's address as its destination IP address.[12] In the example above, the user's e-mail will cross an international link on its way to the foreign-based VPN server, and again on its way from the VPN server to the U.S. recipient. Similarly, if persons located in the United States use foreign-based servers (perhaps belonging to an overseas corporation, or university) to send and receive e-mail from other persons in the United States, those communications will cross international links as they are routed to and from the foreign servers.

58. NSA acquisition of such communications, assuming it occurs as the FISC described, would be entirely consistent with white- or blacklisting by IP address. In the above examples, on the first leg from the sender to the foreign VPN or e-mail server, the destination IP address of the communications will be the foreign IP address of the server; and on the second leg to the recipient the source IP address of the communications will also be the foreign IP address of the server. Therefore, if the NSA were whitelisting communications to and from certain IP addresses, communications of the kind described above could still be copied and scanned by the NSA (and acquired if they contain targeted selectors), if the IP addresses of the foreign VPN and e-mail servers were included on the whitelist. Equally so, if the NSA were blacklisting communications to and from certain IP addresses, wholly domestic communications of the kinds described above could still be copied and scanned by the NSA (and acquired if they contain targeted selectors), if the IP addresses of the foreign VPN and e-mail servers were not on the NSA's blacklist.

59. U.K. Section 8(4) collection: Mr. Bradner states that his conclusion that the NSA likely follows his copy-all-then-scan approach is "reinforced" by public filings of the U.K. Government in the European Court of Human Rights (ECHR) concerning its "Section 8(4)" collection program conducted by the U.K.'s Government Communications Headquarters (GCHQ). Bradner Decl. ¶¶ 368-69. I see little support, if any, for Mr. Bradner's conclusions in these non-

[12] Once the packets reach the VPN server, the VPN layer of encryption is removed, and the packets are forwarded to their intended destination, but for security are assigned the VPN server's address as their source IP address. The process is reversed for any response to the user's initial communication.

technical documents, which contain information, not discussed in his declaration, that tends to refute, not support, his views.

60. The mere fact that the U.K. Government conducts Section 8(4) surveillance using one possible configuration does not mean that the NSA conducts Upstream surveillance the same way, as opposed to other possible approaches that could be followed. Beyond that observation, the description of Section 8(4) collection that Mr. Bradner relies on is contained in publicly filed legal briefs and a court opinion, not internal technical or operational manuals or schematics detailing the design and operation of Section 8(4) collection systems. *See* Bradner Decl. App'x DD, EE. These sources provide only the roughest outline of the Section 8(4) collection process, and it is difficult, therefore, to draw detailed technical conclusions about the Section 8(4) process or how it compares to Upstream collection based on these sources.

61. To the extent that documents at such a high level of generality can be relied upon for the purpose of drawing conclusions about how Section 8(4) collection (or Upstream) operates, I note that they actually describe 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. In a passage not cited in Mr. Bradner's declaration, the U.K. Government's brief before the ECHR describes Section 8(4) collection as follows:

> First stage: collection
>
> GCHQ selects which bearers [circuits] to access based on an assessment of the likely intelligence value of the communications they are carrying. . . .
>
> Second stage: filtering
>
> GCHQ's processing systems operate on the bearers which it has chosen to access. A degree of filtering is then applied to the traffic on these bearers, designed to select communications of potential intelligence value. As a result of this filtering stage, the processing systems automatically discard a significant proportion of the communications on the targeted bearers.
>
> Third stage: selection for examination
>
> The remaining communications are then subjected to the application of queries, both simple and complex, to draw out communications of intelligence value. Examples of a simple query are searches against a "strong selector" such as a telephone number or email address. Complex queries combine a number of criteria, which may include weaker selectors but which in combination aim to reduce the odds of a false positive. Communications that do not match the chosen

> criteria are automatically discarded. The retained communications are available to analysts for possible examination.

Bradner Decl., App'x EE at 4-5.

62. To summarize, according to the above description of Section 8(4) collection in the U.K. Government's brief, before communications are "querie[d]" for the presence of "selectors" (at the third stage), the Section 8(4) processing systems apply "filtering" (at the second stage) to winnow communications deemed to lack intelligence value and to pass on to stage three only those communications considered to be "of potential intelligence value." The exact type of filtering performed at the second stage is not disclosed, but the general description of the Section 8(4) process contained in these publicly available documents is consistent with application of the filter-then-copy-and-scan techniques I have described.

63. Mr. Bradner focuses attention on the U.K. Government's explanation that in order to conduct these filtering and querying processes it is necessary as a "practical" matter, for "technical reasons . . . to intercept the entire contents of a [circuit]." Bradner Decl. ¶ 368 (citing App'x EE ¶¶ 7-8, at 2-3). These "technical reasons" are not described, so there is no way to know whether they would constrain the NSA's ability to configure its collection systems.

64. Assuming as does Mr. Bradner that the "intercept[ion]" referred to in the U.K. brief involves duplication of the entire communications stream before communications are filtered, it does not necessarily follow that the U.K. Government or the NSA must be given access to copies of all communications traveling across a monitored link. If a provider does not prefer to use its network router or switch to perform the IP address or port or protocol number filtering, then it would also be technically feasible for a provider to use an optical splitter, as both Mr. Bradner and I have discussed, to duplicate the communications stream and divert the copied stream for off-line processing while the communications in the "original" stream continue toward their intended destinations. *See* First Decl. ¶ 55; Bradner Decl. ¶¶ 275-76. The provider then could apply IP address or port or protocol number filtering (whether whitelisting or blacklisting)

to the copied stream using a router or standard firewall "appliance"[13] to make available to GCHQ (or the NSA) only those of the copied communications meeting the filter criteria, while automatically destroying the rest. Either configuration would be consistent with what is said in the U.K. Government's brief concerning the "interception" of all communications on a circuit, and neither would involve, much less require, passing all the communications on a monitored circuit to the GCHQ's or the NSA's possession and control.

65. In short, the documents cited by Mr. Bradner offer no basis for concluding (as opposed to speculating) that in conducting Upstream collection the NSA copies and scans all communications, or even that it acquires copies (which it then scans) of all communications, that cross a monitored Internet backbone link.

66. The EINSTEIN 2.0 System: Wikimedia also states in its brief that Mr. Bradner's conclusions are "corroborated" by the U.S. Government's cyber-defense system known as EINSTEIN 2.0. Wikimedia Brief at 20. In contrast, while Mr. Bradner refers to the EINSTEIN 2.0 system as an example of a deep-packet-inspection system, he does not cite EINSTEIN 2.0 as corroboration for his conclusions. Bradner Decl. ¶ 259.

67. EINSTEIN 2.0 is not, as Wikimedia describes it, a "surveillance program" like Upstream collection. Rather, as described in the Department of Justice Office of Legal Counsel ("OLC") memorandum relied on by Wikimedia, Plaintiff's Exhibit 25, EINSTEIN 2.0 is a cyber-intrusion detection system meant to protect the unclassified information technology systems of civilian U.S. Government agencies against malware and other network-based attacks. It is not meant for intelligence gathering, except possibly tracing the progression of detected cyber-attacks.

68. To perform its function, EINSTEIN 2.0 scans incoming Internet traffic as it reaches the access points connecting these Federal Government systems to the Internet. Plaintiff's Exhibit 25 at 3. As described in the OLC memorandum, "EINSTEIN 2.0 sensors [do] not scan actual

[13] The Barracuda CloudGen firewall is one commercial example, with throughput of up to 46 Gb/s.

Federal Systems Internet Traffic for malicious computer code as that traffic is in transmission, but instead will scan a temporary copy of that traffic created solely for the purpose of scanning by the sensors," while "[t]he 'original' Federal Systems Internet Traffic will continue to its destination without being scanned." Plaintiff's Exhibit 25 at 4. There are at least two reasons why conclusions about the Upstream collection process cannot be drawn from this statement about the configuration of the EINSTEIN 2.0 system.

69. First, because 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. For the reasons I have discussed above, and in my first declaration, the NSA could reliably obtain access to its targets' communications crossing a monitored link without copying and scanning all communications that cross that link. Second, it is unlikely that the volume of incoming Internet traffic at any given civilian Government agency exceeds 10 gigabits per second.[14] In comparison, the potential volume of traffic at the AEConnect link, discussed above, 40 terabits per second, is over 4,000 times greater. The relatively small volume of traffic that EINSTEIN 2.0 can be expected to support is well within the capability envelope of a single commercial intrusion detection system or DPI system, meaning that there are likely no processing or capacity constraints in the EINSTEIN 2.0 system that would necessitate filtering out communications to reduce the volume that needs to be scanned for malicious code.

70. Because the purpose and required processing capacity of EINSTEIN 2.0 differ fundamentally from those of Upstream collection, the architecture and operation of EINSTEIN 2.0 are unlikely to provide insight into the operational practices of the Upstream program.

71. Comprehensiveness: Finally, I address Mr. Bradner's and Wikimedia's attempts to draw support from what Mr. Bradner describes as "the NSA's stated desire to be comprehensive

[14] The GSA Enterprise Infrastructure Solutions (EIS) guide supports this conclusion, as the highest available "[d]edicated burstable Internet bandwidth" is 10 Gb/s, as CLIN 22006. (https://eis-public-pricer.nhc.noblis.org/ajax.php/resources/download?type=csv&file=clins) Similarly, the MTIPS (Managed Trusted Internet Protocol Services) item that incorporates EINSTEIN functionality tops out at 10 Gb/s (same file, CLIN MT00060). See also https://www.gsa.gov/technology/technology-products-services/it-security/trusted-internet-connections-tics.

in its [Upstream] collection." Bradner Decl. ¶¶ 228, 333, 359 (citing PCLOB Section 702 Report at 10, 123, 143); Wikimedia Brief at 21. Mr. Bradner infers that "if the NSA's goal is to comprehensively obtain its targets' communications, then it must comprehensively copy, reassemble and review all transactions that could conceivably be to or from a target that transit the circuits being monitored." Bradner Decl. ¶ 335. There are numerous reasons why this conclusion does not follow.

72. The "repeated[ ]" statements by the "[G]overnment" that Mr. Bradner refers to, Bradner Decl. ¶ 333, are actually a single statement by the PCLOB that appears twice in the PCLOB's Section 702 Report. *See, e.g.,* Bradner Decl. ¶ 333 & n.106 (citing PCLOB Section 702 Report at 10, 123). In the cited statement, the PCLOB characterizes "the NSA's acquisition of 'about' communications" as "an inevitable byproduct of the government's efforts to comprehensively acquire communications that are sent to or from its targets." PCLOB Section 702 Report at 10; *see id.* at 123 (same).

73. We cannot simply indulge an assumption that the NSA "comprehensively" acquires the communications of its targets based on the slim reed relied on by Mr. Bradner. I agree with Mr. Bradner that it would be "unsurprising" to discover that the NSA, in a perfect world, would prefer to obtain access to all of its foreign-intelligence targets' communications. It does not follow, however, that the NSA is in fact doing so. It is one thing to state these goals, and quite another to design, construct, deploy, maintain, and pay for the collection systems required, in the numbers and with the capacity needed, to attain such ambitious goals. We cannot assume on the basis of a stated goal alone that the NSA has achieved that desired result without assuming away the technical, logistical, and financial hurdles, the resource constraints and trade-offs, and the competing mission priorities, that would stand in the way. Even if the technical and logistical hurdles could be overcome, we must recognize the possibility that at some point the cost of doing so may, in the NSA's view, outweigh the marginal benefit of potentially discovering still further communications of its targets in some as-yet unexplored stream of communications on the Internet. In short, we must recognize that in the field of large-

scale digital communications engineering as in other practical endeavors, idealized goals must be tempered by gritty reality.

74. Even taking for granted, for a moment, Mr. Bradner's assumption that the NSA has achieved the goal of comprehensively acquiring its targets' online communications, it still would not follow that the NSA "must comprehensively copy, reassemble and review all [communications] . . . transit[ing] the circuits being monitored." Bradner Decl. ¶ 335. I explained in my first declaration how the NSA, using traffic-mirroring techniques such as white- and blacklisting, could reliably obtain access to its Upstream targets' communications without copying and scanning all of the communications traversing a monitored link. And, as I explain herein, Mr. Bradner identifies no technical reason to question that conclusion, instead relying on speculation about the nature and habits of the NSA's targets, the NSA's intelligence priorities, its resources, and capabilities, to support his opinions to the contrary.[15] Moreover, it must be acknowledged that the term comprehensive is a qualitative one, susceptible of a number of meanings other than "exhaustive." Without detailed information, at the least, concerning the types and quantities of communications accessed by the NSA, it is not possible to reverse-engineer detailed conclusions about the methodologies employed and equipment configurations deployed by the NSA from so limited and technically unenlightening a starting point as the supposed comprehensiveness of the agency's objectives.

75. To reiterate, I do not opine on the likelihood that the NSA, in the course of conducting Upstream surveillance, actually may use the traffic-mirroring techniques I have discussed. But it is my opinion that, at bottom, the reasons given by Mr. Bradner for deeming

[15] Mr. Bradner also relies on the PCLOB's characterization of Upstream's objective to conclude that the NSA "is very likely to be monitoring a large number of international circuits, given that it would need to monitor most, if not all, such circuits to accomplish its stated (and unsurprising) goal of reliably and comprehensively collecting the communications of its targets." Bradner Decl. ¶ 353. As I have stated, one cannot simply take for granted that the PCLOB's 2014 description of the NSA's supposed goals reflects the reality of what the NSA has actually accomplished. Be that as it may, I have also already explained that the feasibility of whitelisting and blacklisting does not depend on the number of sites on the Internet that the NSA actually monitors, whether the number is one or many. First Decl. ¶¶ 90-91. Mr. Bradner does not contest this point.

that possibility "implausible," and his alternative approach "most likely," are without basis in Internet technology and engineering.

MR. BRADNER'S CERTAINTY THAT THE NSA HAS COPIED AND SCANNED AT LEAST SOME WIKIMEDIA COMMUNICATIONS, EVEN IF THE NSA EMPLOYS TRAFFIC-MIRRORING TECHNIQUES, IS WITHOUT TECHNICAL OR EMPIRICAL BASIS.

76. Although Mr. Bradner acknowledges that the various types of traffic-mirroring techniques I discuss are technically feasible, he nevertheless maintains that, even if the NSA were using these techniques to filter out Wikimedia communications from those made available to it at a monitored link (whether by design, or effect), it is still "virtually certain" that the NSA, in the course of Upstream collection, "has copied, reassembled, and reviewed at least some of Wikimedia's communications." Bradner Decl. ¶ 370. He does not explain, however, why he believes this to be the case. Although he raises a number of reasons why he believes the NSA likely would not employ these techniques (which I addressed in the preceding sections), he raises no technical objections to the *efficacy* of white- or blacklisting port or protocol numbers, or whitelisting by IP address, as means of blocking access to Wikimedia communications.

77. Rather, Mr. Bradner asserts only that it is "technologically incorrect that blocking [i.e., blacklisting] Wikimedia IP addresses would block all Wikimedia traffic." Bradner Decl. ¶ 367(b). So far as I can discern from Mr. Bradner's declaration, this is the sole "technical[ ] inaccura[cy]," Bradner Decl. ¶ 7(b), that Mr. Bradner ascribes to my analysis. Specifically, Mr. Bradner posits three scenarios in which Wikimedia's communications "would still be copied, reassembled and reviewed by the NSA," even if the NSA blacklisted communications containing Wikimedia IP addresses. Bradner Decl. ¶ 367(b). Each of these scenarios, while theoretically possible, could come to pass only in the uncertain event that particular conditions are met, as I discuss in turn.

78. Scenario One: MCT Containing a Wikimedia Communication: Mr. Bradner observes that if a Wikimedia communication were contained within a so-called "multi-communication transaction" (MCT) that itself was neither to nor from Wikimedia, then blacklisting communications with source or destination IP addresses assigned to Wikimedia

would not prevent the MCT, including the embedded Wikimedia communication, from being copied and scanned at a monitored link. Bradner Decl. ¶ 367(b)(1). (This would be so, because the IP addresses included in the individual packet headers would be the source and destination IP addresses of the enclosing MCT, rather than the IP addresses of the embedded Wikimedia communication.) Mr. Bradner gives, as a hypothetical example of an MCT, the group of e-mail messages that are transmitted together as a single communication from an e-mail service to a subscribing user's Inbox, when the user logs in to check his or her e-mail. Bradner Decl. ¶¶ 132, 317.[16] This situation would result in copying and scanning communications to or from Wikimedia, however, only if several conditions were met:

a. First, Wikimedia maintains that its communications are copied and scanned as they traverse international links on the Internet backbone. For an MCT comprised of e-mails downloaded from a server to the Inbox on an individual user's cellphone or computer to cross an international Internet link, either a user in a foreign location must be downloading e-mails from a server located inside the United States, or a user located in the United States must be downloading e-mails from a foreign server.

b. Second, for the MCT containing these e-mails to include Wikimedia communications, the user must also be one who communicates by e-mail with Wikimedia.

c. Third, for the MCT to be copied and scanned by the NSA, it would have to traverse an international link that is monitored by the NSA.

d. Fourth, the MCT in which the Wikimedia communication is embedded must itself be a communication that has not been blacklisted.

79. Wikimedia claims to communicate with persons in almost every country on Earth, and that it is "virtually certain," therefore, that its communications traverse every circuit on every international cable carrying Internet traffic to and from the United States, *see* Bradner Decl. ¶ 6(d). But the scenario posited by Mr. Bradner is limited to those persons *meeting the first two criteria above.* Neither Wikimedia nor Mr. Bradner cites evidence concerning the number or geographic locations of persons meeting those criteria, if any, who communicate with Wikimedia, and I am aware of none. There is no basis, therefore, to conclude that MCTs enclosing Wikimedia communications almost certainly cross every international Internet link to and from the United

[16] The NSA has not publicly acknowledged what kind of MCTs are acquired during the Upstream collection process.

States, or, critically, that they cross one or more links that happen to be monitored (assuming any) by the NSA (the third criterion). It is a matter of conjecture, therefore, whether MCTs containing Wikimedia communications are copied and scanned by the NSA while crossing international Internet links.

80. The scenario is rendered even more conjectural when the possibility is considered that an MCT in which a Wikimedia communication is embedded might itself be a blacklisted communication (thus failing to meet the fourth criterion). U.S. consumer-based webmail services, such as a person located outside the United States might use, most commonly interface with their subscribers using encrypted HTTPS communications. (Gmail, the most popular webmail service provider, has been 100-percent encrypted since mid-2014.)[17] In addition, most businesses and other organizations strongly discourage or prevent use of unencrypted protocols to send or receive e-mail for business purposes. Thus, for example, employees of U.S.-based companies traveling or stationed abroad who rely on an e-mail server located at a U.S. corporate headquarters most likely will send and retrieve e-mail, respectively, through an encrypted SMTP submission transport connection, using TCP port 587, and an encrypted IMAP connection, using TCP port 993. Overall, the likelihood that a person located outside the United States would use unencrypted e-mail protocols to send or receive e-mail is exceedingly small. As I observed in my earlier declaration, if the NSA lacked the ability to decipher a particular kind (or kinds) of encrypted communications, then it could avoid copying and scanning them, if it wished, by blacklisting their associated port or protocol numbers. First Decl. ¶ 79. That would include encrypted MCTs that might hypothetically contain Wikimedia communications.

81. Scenario Two: E-mail to Wikimedia from Abroad Using a U.S.-Based Service: Mr. Bradner next posits that the NSA could copy and scan Wikimedia communications, even if it blacklisted Wikimedia's IP addresses, in a "case where a person located outside the U.S. is using an email service located inside the U.S. to send email to [and receive email from] Wikimedia." Bradner Decl. ¶ 367(b)(2). (This would be possible because while in transmission from the user

[17] https://transparencyreport.google.com/https/overview.

to the e-mail server, and vice versa, the communication would not include Wikimedia's IP address in the packet headers.) In practical terms, this scenario is simply a variant of the first, and could occur only if the following conditions, quite similar to those in the MCT scenario above, are met:

a. First, for the e-mail in question to cross an international Internet link without a Wikimedia IP address in the packet header, the user must be someone located outside the United States who is using an e-mail service (more precisely, an e-mail server) that is located inside the United States.

b. Second, the user must be sending e-mail to and/or receiving e-mail from Wikimedia.

c. Third, for the e-mail to be copied and scanned by the NSA, it would have to traverse an international link that is monitored by the NSA.

d. Fourth, the e-mail must itself be a communication that has not been blacklisted.

82. As in Mr. Bradner's MCT scenario, there is no basis on which to conclude that e-mail between Wikimedia and users who meet the first two criteria above (if any) almost certainly cross every international Internet link to and from the United States, or that they cross one or more links that happen to be monitored (assuming any) by the NSA (the third criterion). It is also a matter of conjecture, therefore, whether e-mail sent to Wikimedia, even from persons outside the United States who are using U.S.-based e-mail services, are copied and scanned by the NSA while crossing international Internet links. And, as in the MCT scenario, the matter becomes even more uncertain when one considers that any international e-mail communications of the kind posited by Mr. Bradner are quite likely encrypted with secure e-mail and/or HTTPS transmission protocols, raising the possibility that they are blacklisted (thus not meeting the fourth criterion), and so excluded from copying and scanning.

83. Scenario Three: Accessing Wikimedia Websites from Abroad Using a U.S.-Based VPN Service: Finally, Mr. Bradner posits that the NSA could copy and scan communications to and from a Wikimedia website, even if it blacklisted Wikimedia's IP addresses, if a user outside the United States accessed the Wikimedia site using a VPN (virtual private network) service located inside the United States. Bradner Decl. ¶ 367(b)(3). This is possible because, as discussed in paragraph 57, above, when a user communicates via a VPN, all of the user's communications

(including HTTP and HTTPS communications) are first routed through the VPN server, and assigned the server's address as the communications' destination IP address, rather than (as in this scenario) the target website. In reality, however, the scenario envisioned by Mr. Bradner could occur only if the following conditions are met:

a. First, for the HTTP and HTTPS communications in question to cross an international Internet link, the user must be someone located outside the United States.

b. Second, the user must be someone who has decided to use, and perhaps to pay a service fee for, the VPN service. (The user may also be using the communication facilities of an organization, such as an employer, that has decided to use a VPN for its business communications.)

c. Third, although the user is located outside the United States, the user (or the organization whose facilities he or she is using) must have chosen a VPN service based inside, not outside, the United States.

d. Fourth, to be copied and scanned by the NSA, the user's communications with Wikimedia's websites would have to traverse an international link that is monitored by the NSA.

e. Fifth, communications to the VPN server must not themselves be blacklisted.

84. Again, there is no basis on which to conclude that communications with Wikimedia websites from users who meet the first three criteria above (if any) almost certainly cross every international Internet link to and from the United States, or that they cross one or more links that happen to be monitored (assuming any) by the NSA (the fourth criterion). It is also a matter of conjecture, therefore, whether HTTP or HTTPS communications sent to Wikimedia websites, even those from persons outside the United States who are using U.S.-based VPN services, are copied and scanned by the NSA while crossing international Internet links. And it becomes even less certain considering that any such communications would be encrypted by the VPN service, raising the possibility, if the NSA lacked the ability to decipher them, that they are blacklisted (perhaps using the VPN's IP address), and thus, not meeting the fifth criterion, are excluded from copying and scanning.

85. In sum, it is far short of the certainty claimed by Mr. Bradner that any of the scenarios described by him would come to pass at a particular international Internet link that happened to be monitored by the NSA (if any), such that the NSA would copy and scan

communications of Wikimedia's even if it had blacklisted Wikimedia IP addresses. Each of these scenarios requires a confluence of multiple events before it could come to pass, the likelihood of which individually, and certainly collectively, is at best conjectural. And so, while I acknowledge that in theory blacklisting might not eliminate all possibility that the NSA, in conducting Upstream surveillance, obtains copies of and scans Wikimedia communications for targeted selectors, blacklisting would render that possibility a matter of speculation.

86. Moreover, it bears emphasis that even if there remained a possibility of copying and scanning Wikimedia communications despite blacklisting Wikimedia IP addresses, that would not be the case if the NSA were employing a whitelisting technique. In a whitelisting scenario, no communications to or from Wikimedia would be copied and made available for scanning by the NSA, unless a communication to or from Wikimedia itself had a source or destination IP address, respectively, on the target whitelist. Nor does Mr. Bradner suggest otherwise. Therefore, the hypothetical possibility that Wikimedia communications could be copied and reviewed in the limited and uncertain sets of circumstances suggested by Mr. Bradner, notwithstanding blacklisting, does not alter my conclusion that the NSA, through whitelisting, could conduct Upstream-type surveillance as envisioned by Wikimedia, without copying and reviewing or otherwise interacting with Wikimedia communications.

87. For these reasons, I find Mr. Bradner's assertion that it is "virtually certain" the NSA has copied and scanned at least some of Wikimedia's communications, even if the NSA employs one or more of the whitelisting and blacklisting techniques I have described, Bradner Decl. ¶ 370, to be without a basis in Internet technology and engineering that rises above the level of conjecture.

PRACTICAL REASONS FOR MAINTAINING "HTTPS BY DEFAULT" AND IPSec ENCRYPTION IN THE CURRENT COMMUNICATIONS ENVIRONMENT

88. Finally, I discuss (a) the significant reasons why, in the current digital communications environment, any organization that operates a major website would be powerfully motivated to protect communications to and from its site using the HTTPS protocol,

in particular HTTPS by default, and (b) the reasons why an organization that transmits sensitive proprietary or personal information on the Internet would be equally motivated to encrypt those communications using a security protocol such as IPsec or ssh.

89. As I explained in my earlier declaration, communications on the Internet may be encrypted to protect the privacy and integrity of the information they contain. The HTTPS protocol (technically, a combination of the Transport Layer Security (TLS) protocol and the HTTP protocol), is the most common encryption mechanism on the Internet, used to ensure that a user's browser connects to the correct web server (rather than an imposter site), and that the information sent between the user and the website can be read only by the user's web browser and the host web server, but not third parties. First Decl. ¶ 42.

90. Although selective encryption of web communications for online transactions deemed sensitive or confidential (such as online banking) began as long ago as 1994, in 2014 the Internet Engineering Task Force (IETF) (*see* First Decl. ¶ 26) described a range of motivations for using encryption pervasively in today's environment, including: (i) surveillance by nation-state actors; (ii) legal but "privacy-unfriendly" exploitation of user information by commercial enterprises; and (iii) various forms of cybercrime.[18] Online surveillance by nation-state actors is a global phenomenon, not limited to activities conducted by the NSA or even the U.S. Government.[19] Large Internet service providers have indicated their interest in aggregating and monetizing data about the web-browsing patterns of their subscribers by selling data to online advertisers.[20] For example, in March 2016 the Federal Communications Commission imposed a $1.35 million fine on Verizon Wireless for its use of a technology that allowed marketers to track customers' web browsing, without their knowledge, so they could be provided more targeted

[18] IETF RFC 7258 (May 2014), https://tools.ietf.org/html/rfc7258

[19] *See* https://rsf.org/en/news/special-report-internet-surveillance-focusing-5-governments-and-5-companies-enemies-internet; https://en.wikipedia.org/wiki/List_of_government_mass_surveillance_projects.

[20] *See* https://arstechnica.com/tech-policy/2017/03/ad-industry-lobbyists-celebrate-impending-death-of-online-privacy-rules/

online advertising.[21] Cybercrime, such as "hacking" users' credit card numbers or login credentials for purposes of fraud, or theft, is widespread on the Internet.[22] Commenting on this environment in an October 2015 address, Wikipedia founder Jimmy Wales reportedly stated, "There's really no excuse to have any major web property that's not secure[d]" using HTTPS.[23]

91. Although for nearly two decades encryption was used primarily to facilitate online banking transactions and credit card purchases, over the last five years HTTPS encryption of web communications has increasingly become the norm, and a security "best practice." Many websites now automatically redirect visitors to the secure (HTTPS) versions of their sites even if a user's browser first contacted the unencrypted (HTTP) version. By 2017, half of all traffic was encrypted, according to Mozilla, developer of the Firefox browser[24], and over 70 of the top 100 most popular websites worldwide had enabled HTTPS encryption, up from 37 in 2016, according to reports by Google.[25] Google reports that by January 2019 the number had risen to 96 of the 100 top sites (which account for approximately 25 percent of global web traffic), and that as of February 2019 over 90 percent of the webpages viewed using Google Chrome (the world's most popular web browser) were loaded over HTTPS connections.[26] CloudFlare, one of the largest content distribution networks (used to facilitate worldwide distribution of client websites' content to end users) now offers redirection to HTTPS encryption as a standard feature of its services.[27]

92. Any organization that continues to operate a website in the current digital environment without at least offering, if not requiring, a secure HTTPS encryption places at risk

---

[21] *See* https://www.cnet.com/news/verizon-racks-up-1-35-m-bill-for-violating-consumer-privacy/.

[22] *See* https://www.pcworld.com/article/205051/Norton_Study_Says_Cybercrime_is_Rampant.html.

[23] https://motherboard.vice.com/en_us/article/ezvj8k/jimmy-wales-theres-really-no-excuse-not-to-use-encryption

[24] *See* https://www.wired.com/2017/01/half-web-now-encrypted-makes-everyone-safer/

[25] *See* https://www.zdnet.com/article/google-this-surge-in-chrome-https-traffic-shows-how-much-safer-you-now-are-online/.

[26] *See* https://transparencyreport.google.com/https/overview?hl=en.

[27] *See* https://www.cloudflare.com/lp/ssl-for-saas-providers/.

the online privacy and security of its users, and perhaps its own proprietary information and commercial interests. Online vendors whose operation depends on the willingness of users to visit their sites and share sensitive personal information (like credit card numbers) to make purchases would discourage visitors and lose business if they did not offer a secure connection. Non-commercial sites that deal with sensitive subject matters, or that simply value online privacy as a matter of principle, also face increasing pressure to implement HTTPS encryption to protect the confidentiality of their users' communications.

93. The transition is being spearheaded through efforts by Google, and other online organizations, to promote the adoption of HTTPS encryption across the entire web. Since July 2018, Google's Chrome web browser has been labeling all websites that do not use the HTTPS protocol as "Not Secure," a warning that may be interpreted by visitors to mean that the site has been hacked or compromised, or may even be malicious. Mozilla's Firefox web browser is reportedly following suit.[28] Given the increasing demands today for greater Internet security, labeling a website as "not secure" could bring reputational damage to the site (and perhaps its parent organization), and would be contrary to the interests of any organization seeking to maximize visits to its site, whether its content is considered sensitive or not.

94. At this time two methods of HTTPS encryption are available. Opt-in HTTPS gives a user the option, when he or she visits a website, of communicating over an HTTP connection or choosing a secure HTTPS connection. In contrast, HTTPS by default involves the automatic redirection of a user to the HTTPS version of the website, even if the user's browser first seeks to make an unsecure HTTP connection. HTTPS by default also involves use of the HSTS (HTTP Strict Transfer Security) protocol to reconfigure the user's browser to connect automatically via HTTPS transmission whenever the user again visits the same site, and to refuse connection to the unencrypted site. Both versions require roughly the same initial and ongoing investment of resources to develop, operate, and maintain the required technical infrastructure, including the

[28] *See* https://www.deepdotweb.com/2018/01/05/mozilla-label-http-sites-not-secure-future-versions-firefox/

retention of qualified engineering personnel. In some circumstances, opt-in HTTPS encryption may be marginally less costly for an organization to implement, depending on the fraction of unencrypted HTTP traffic visiting the website, because it may require somewhat less server capacity than supporting all site visitors exclusively via HTTPS.

95. All other things being equal, the level of encryption protection provided is the same whether a web communication is encrypted with HTTPS optionally or by default. The difference between the two methods is that HTTPS by default adds a level of protection against so-called SSL stripping attacks. In an SSL stripping attack, if a user's browser makes an unsecure HTTP request to make a secure HTTPS connection to a website, an attacker can intercept that initial, unsecure request and use it to "strip" the encryption that otherwise would have protected the ensuing exchange of communications between the user and the website. The attacker thus obtains access to the contents of the user's communications with the site, including the information the user views or downloads from the site, and any sensitive personal or financial information the user shares with the site. Implementing HTTPS by default mitigates the threat of such attacks through the HSTS reconfiguration of a user's browser to encrypt future connection requests after the user's first visit to the site, or if the site is listed on the HSTS preload list incorporated into the user's browser. Largely for this reason, the adoption of default rather than opt-in HTTPS, like the adoption of HTTPS encryption generally, is increasingly becoming the norm, particularly by sites where users must enter login credentials or provide sensitive personal information in order to conduct financial transactions. For example, nine out of 11 U.S. banking sites and all 13 U.S. Government sites listed at the HTTPSWatch website support HSTS, and all but one redirect from HTTP to HTTPS automatically[29]. This confirms the prediction of an August 2017 study[30]: "Top websites will be almost entirely HTTPS within a year and a half. Half have

[29] See https://httpswatch.com/us; visited Feb. 9, 2019.

[30] Adrienne P. Felt, et al., "Measuring HTTPS Adoption on the Web," 26th USENIX Security Symposium, August 2017, https://www.usenix.org/system/files/conference/usenixsecurity17/sec17-felt.pdf.

moved, more are preparing to move, and the rem[a]inder will feel pressured to meet the changing industry standard."

96. So far as I am aware, SSL stripping attacks have never been publicly identified as a surveillance technique employed by the NSA in connection with its Upstream collection program, whether in any official U.S. Government disclosures. Such an attack would require a man-in-the-middle attack or DNS redirection, not just passive intercept.

97. For the same wide variety of reasons identified by the IETF, going well beyond what has been publicly reported about NSA Upstream surveillance, encryption increasingly has become the default across the Internet as a whole, not simply on the World Wide Web. IPsec, the Internet Protocol Security suite, is a set of network security protocols commonly used by commercial enterprises and other large organizations to encrypt sensitive data that they transmit from one business site to another, such as between large data centers. Wikimedia has implemented IPsec both for the transmission of its web server logs between its foreign servers in the Netherlands and its servers located in the United States, *see* First Decl.¶ 83; Gov't Exh. 4 (Technical Statistics Chart), and apparently for the transmission of cache-to-cache data between its U.S. data centers, Plaintiff's Exhibit 3 (Paulson Decl.) ¶ 53; Plaintiff's Exhibit 39 at WIKI00006566. The transmission of a website's server logs in encrypted form is an accepted best practice, not only to protect proprietary information about the operation of the site, but to protect user information that may be considered personal. Indeed, the European Union's General Data Protection Regulation, which became effective in May 2018, classifies information contained in web server logs (principally IP addresses) as personal data that must be handled in a manner that ensures appropriate security, such as by encrypting them.[31] Finally, other internal communications that traverse the Internet, such as cache-to-cache communications, are protected by encryption primarily to ensure integrity, e.g., to prevent malicious actors from modifying web pages in transit. While this threat is less likely than others, ISPs in China, Russia and Pakistan, among other countries, have temporarily diverted traffic through their country. If

[31] https://www.ctrl.blog/entry/gdpr-web-server-logs

a third party could modify web pages in transit, they could install trackers or malware, for example.

98. Even if the NSA were not conducting Upstream surveillance, there would remain numerous reasons, discussed above, why an organization would be highly motivated to encrypt transmissions of its web server logs and cache-to-cache data. The online transmission of such sensitive proprietary or private information using unencrypted protocols not only poses known risks of interception, modification and theft by unauthorized third parties (including foreign government actors), it has also become an indicator that an organization lacks proper cyber hygiene.

CONCLUSION

99. For the reasons I discuss above and in my first declaration, it remains my opinion that, based on what is publicly known about the NSA's Upstream collection technique, the NSA in theory could be conducting this activity, at least as Wikimedia conceives of it, in a number of technically feasible, readily implemented ways that could avoid NSA interaction with Wikimedia's online communications.

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

101. It is also my opinion that even if the NSA were not conducting Upstream surveillance, in the current digital communications environment there would remain numerous reasons for an organization that operates one or more major websites to implement HTTPS-by-

default on its websites, and for an organization that transmits large volumes of proprietary or other sensitive data across the Internet to encrypt those transmissions using IPsec or another such encrypted transmission protocol.

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

Executed in New York, New York on February 15, 2019.

HENNING G. SCHULZRINNE