CASE NO. 20-1191

UNITED STATES COURT OF APPEALS
FOR THE FOURTH CIRCUIT

WIKIMEDIA FOUNDATION,

Plaintiffs and Appellants,

vs.

NATIONAL SECURITY AGENCY, *et al.*,

Defendants and Appellees.

On Appeal from the United States District Court
for the District of Maryland
Baltimore Division

BRIEF OF *AMICI CURIAE* NETWORK ENGINEERS AND TECHNOLOGISTS IN SUPPORT OF PLAINTIFF-APPELLANT WIKIMEDIA AND REVERSAL

MUNGER, TOLLES & OLSON LLP
Jonathan Blavin (CA Bar 230269)
Elizabeth Kim (CA Bar 295277)
Alexander Gorin (CA Bar 326235)
560 Mission Street
27th Floor
San Francisco, CA 94105
Telephone: (415) 512-4011
Facsimile: (415) 644-6911
jonathan.blavin@mto.com

Attorneys for *Amici Curiae*

TABLE OF CONTENTS

Page

TABLE OF AUTHORITIES

Page(s)

OTHER AUTHORITIES

INTEREST OF AMICI CURIAE

Amici curiae are a group of computer scientists, network engineers, professors, Internet networking experts, and academics with diverse expertise on the science and practice of Internet networking. *Amici* have unique experience to explain and clarify the technological facts underpinning the parties' claims. **All *Amici* sign in their personal capacity, and titles and employer affiliations are provided for identification purposes *only*.** *Amici* include the following individuals as well as the additional signatories identified in Appendix A.

David Crocker worked in the ARPAnet and NSF-CSNet research communities. He then managed product teams and founded several startup companies. Dave was co-recipient of the 2004 IEEE Internet Award for his work on email. Over forty-five years of developing standards, he has authored 65 IETF Requests For Comments, including internet mail, instant messaging, facsimile, EDI, and security.

Bruce Schneier is an internationally renowned security technologist, called a "security guru" by The Economist. Schneier is a fellow at the Berkman Center for Internet and Society at Harvard University, a board member of the Electronic Frontier Foundation, and an Advisory Board member of the Electronic Privacy Information Center. He is also the Chief of Security Architecture at Inrupt, Inc.

Patrick W. Gilmore has been chief network architect of national and international backbones as well as Akamai, the largest CDN on the Internet. He has been CTO of a cloud computing company and recently co-founded a datacenter company. He serves on the boards of NANOG, ARIN, PeeringDB, and Seattle IX, and previously served on the boards of ARIN and LINX.

Accordingly, *amici* respectfully submit this brief in support of appellants.[1]

[1] All parties have consented to the filing of this brief pursuant to FED. R. APP. P. 29(a). No party's counsel authored this brief in whole or in part. No party or party's counsel contributed money that was intended to fund preparing or submitting this brief. No person—other than *amici* or their counsel—contributed money that was intended to fund preparing or submitting the brief.

Introduction

This lawsuit concerns a dispute between the Wikimedia Foundation, a non-profit organization primarily known for hosting Wikipedia, and the National Security Agency (NSA), regarding the legality of the NSA's Upstream surveillance program. Wikimedia alleges that the NSA "has intercepted, copied, and collected Wikimedia's Internet communications pursuant to the Upstream surveillance program" and that the collection "exceeds the NSA's authority under FISA" and violates Wikimedia's constitutional rights. *See* SJ Op. at 1-2 (JA 7:4073-74). The government disagrees, stating Wikimedia lacks standing because there are "technically feasible" and "readily implemented" means to selectively copy communications while excluding all of Wikimedia's. 3d Schulzrinne Decl. ¶ 2 (JA 7:4021).

The district court reviewed multiple declarations from the government's expert, Dr. Henning Schulzrinne, discussing a hypothetical technique by which the government could "conduct Upstream surveillance on an international Internet circuit 'without intercepting, copying, reviewing, or otherwise interacting with [the] communications of Wikimedia.'" *See* SJ Op. at 27 (JA 7:4099). Under this theory, the government could use traffic mirroring techniques combined with asking the network operator to filter traffic in a way which would exclude certain communications from being intercepted, copied, and reviewed, and the government

would never be in possession of Wikimedia's communications. 3d Schulzrinne Decl. ¶¶ 5, 18, 20, 26 (JA 7:4022, 4027, 4028, 4030). The district court accepted that "Dr. Schulzrinne has convincingly demonstrated that there are technologically feasible methods by which the NSA could hypothetically operate Upstream surveillance that would result in the NSA not copying or intercepting any of Wikimedia's communications." SJ Op. at 38 (JA 7:4110). The district court credited Dr. Schulzrinne's theories and concluded that copying or collecting of Wikimedia's communications was neither "certainly impending" nor was there a "substantial risk collection will occur." SJ Op. at 44 (JA 7:4116).

We are a group of networking experts and technologists, including many architects and engineers of the very networks being discussed in this case, who are presenting our technical opinion regarding claims made by Wikimedia and the NSA. We disagree with the district court's conclusions.

The government's theory does not hold up to the practical realities of operating large international computer networks. The hypothetical explanations for how the government *could* avoid Wikimedia's communications fall apart in the transition from academic thought exercise to a practical solution implemented on actual networks. We conclude the district court erred in finding that the NSA could

"*hypothetically*, utilize a process of whitelisting and blacklisting[2] to filter out" Wikimedia's communications "prior to scanning the Internet communications for targeted selectors." SJ Op. at 33 (JA 7:4105) (emphasis in original). We have evaluated Dr. Schulzrinne's hypothetical technique by which the government could theoretically, either passively or actively, avoid the totality of Wikimedia's communications using mirror ports combined with whitelist and/or blacklist filters. We conclude that using traffic mirroring combined with filtering to conduct surveillance as the government describes in its litigation materials lacks a basis in both Internet technology and engineering.

The government's hypothetical must be implemented by the Internet Service Providers ("ISPs"), and would require those networks to perform physically or economically infeasible acts. The hypothetical would also require the NSA to disclose their target list to many ISP employees, many of whom are likely to be non-US citizens. Therefore it would pose a significant risk to the operation of large Internet networks and is neither technically feasible nor readily implementable. For the foregoing reasons, we agree with Wikimedia's expert that it is "virtually certain" that the NSA has copied, intercepted, or reviewed at least some of Wikimedia's

[2] The current terms of art are "allowlist" and "denylist." However, since the previous declarations use the outdated terms "whitelist" and "blacklist," we will follow that convention here to avoid confusion.

communications as a matter of technological necessity. *See, e.g.,* 2d Bradner Decl. ¶¶ 114-15 (JA 7:3919-20).

I. BACKGROUND

A. HOW THE INTERNET WORKS, IN BRIEF

In order to evaluate the hypotheticals proposed by the government, it is necessary to understand a few technical details. The Internet is not a single network, but is more accurately described as a network of networks. Network operators who provide access to the Internet are called "Internet Service Providers" or ISPs. ISPs include well-known providers of broadband service, such as Comcast, Verizon FiOS, and AT&T U-verse; as well as the so-called "tier one" networks such as CenturyLink, Cogent, and NTT. The largest ISPs have links that can transfer enormous amounts of data between cities and across oceans called "backbone" links. These links are connected to routers, which are specialized computers that direct traffic. For this document, we define "Core Router" as a router that connects backbone links, including traffic entering and exiting the United States through international undersea cables.

B. FACTUAL BASIS FOR OUR TECHNICAL ANALYSIS AND CONCLUSIONS

Our conclusions are based solely on information publicly disclosed by the

government, documents filed by the parties in this case[3], the district court opinion, and the signatories' personal knowledge and expertise. We take no position regarding disputes over admissibility of evidence presented by the parties or any other legal disagreements between the parties.

The following are the key facts we used to form our conclusions. First, we credit the district court's conclusion that "the NSA monitors at least one circuit carrying international Internet communications in the course of Upstream surveillance and that Wikimedia's communications traverse every circuit carrying international Internet communications from the United States to the rest of the world." SJ Op. at 37 (JA 7:4109); *see also* Privacy and Civil Liberties Oversight Board, Report on the Surveillance Program Operated Pursuant to Section 702 of the Foreign Intelligence Surveillance Act ("PCLOB Report") 35-37(2014); [*Redacted*], 2011 WL 10945618, at *15 (FISA Ct. Oct. 3, 2011).

We also understand, based on the PCLOB's unclassified report, that the NSA generally aims "to comprehensively acquire communications that are sent to or from its targets." PCLOB Report at 10, 123. The same publicly released report states that a "selector" is a "specific communications facility" such as an "email address or

[3] While some of Appellants' litigation materials contain information from unauthorized leaks of US government information, we did not consider leaked information in forming our conclusions.

telephone number." PCLOB Report at 32-33. We interpret the PCLOB's statement to mean that the NSA collects more than a negligible number of targeted selectors. We also understand the government's hypothetical technique is not designed to exclude solely Wikimedia's communications while pulling in all communications from every other source on the Internet.

Our understanding of the general methodology of the NSA's Upstream surveillance program is based on its description by the district court. According to the district court opinion, Upstream surveillance involves three steps:

> "First, certain Internet transactions transiting the Internet backbone networks of certain electronic communication service providers are filtered for the purpose of excluding wholly domestic communications. Second, these Internet transactions are then scanned to identify for acquisition those transactions that contain communications to or from . . . persons targeted in accordance with the applicable NSA targeting procedures. And third, those transactions that pass through both the filtering and the scanning are ingested into government databases."[4]

SJ Op. at 15 (JA 7:4087) (cleaned up). On top of these public disclosures, the government's expert poses a hypothetical technique by which the government could be avoiding all of Wikimedia's communications. For the reasons discussed in detail

[4] Wikimedia contests the district court's description in its brief, arguing that "the government's disclosures show that it does not perform any filtering when it conducts Upstream surveillance at 'international Internet links.'" Appellant Br. at 11. Our conclusions remain the same whether one accepts the district court's characterization of the program or Wikimedia's characterization.

below, we conclude that the government's argument for how it could theoretically avoid all of Wikimedia's communications is neither "technically feasible" nor "readily implementable." 3d Schulzrinne Decl. ¶ 2 (JA 7:4021).

II. ARGUMENT

A. THE GOVERNMENT MISTAKENLY RELIES ON TRAFFIC MIRRORING AS A POTENTIAL MEANS TO AVOID ALL OF WIKIMEDIA'S COMMUNICATIONS

The government's theory as to how it could avoid all of Wikimedia's communications, as accepted by the district court, is not based in network realities because it problematically relies on the hypothetical use of traffic mirroring to deliberately filter out Wikimedia's communications.

Dr. Schulzrinne states: "There are at least two well-known approaches to obtaining copies of Internet communications at locations other than the sources or destinations of the communications . . . ." 1st Schulzrinne Decl. ¶ 54 (JA 1:0743). He is describing "mirror ports" and "fiber optic splitters." Were the government to use fiber optic splitters to collect data, both the government and Wikimedia agree that the government would copy or intercept at least some of Wikimedia's communications.

Dr. Schulzrinne admits that the use of an optical splitter would require the government to copy all communications on a monitored network, including

Wikimedia's, but dismisses the possibility that the government is using a fiber optic splitter because mirror ports provide a "technically feasible" and "readily implementable" alternative. 3d Schulzrinne Decl. ¶ 2 (JA 7:4021); 1st Schulzrinne Decl. ¶¶ 56 (JA 1:0743-44). Thus, in order to show that Wikimedia lacks standing, the government develops a hypothetical which avoids the use of fiber optic splitters. The district court relies on this "traffic mirroring hypothetical." SJ Op. at 29 (JA 7:4101).

To clarify, a mirror port, or "traffic mirroring," is when a Core Router is configured to copy all traffic from one link to another without interrupting the original copy on its way to the destination. SJ Op. at 12 (JA 7:4084) (briefly describing traffic mirroring). Using whitelists and/or blacklists, it is possible to configure a mirror port to copy and forward some packets, while ignoring others. The use of mirror ports is central to the district court's finding for the defendant. SJ Op. at 29-33 (JA 7:4101-05).

An optical splitter is a device attached to a fiber optic cable carrying electronic communications and reflects a portion of the light traveling down that fiber to a different receiver. In this way, the signal is duplicated, creating an exact copy of the information being transferred. The information continues on its original course to the end user, while an identical copy is sent to the surveilling entity. We agree when

Dr. Schulzrinne states, "[T]he use of fiber-optic splitters to obtain copies of online communications for surveillance purposes would entail, as alleged by Wikimedia, the copying of all communications flowing across a given fiber-optic link." 1st Schulzrinne Decl. ¶ 56 (JA 1:0743-44). This would be the "copy-all-then-scan" approach discussed by the parties' experts. *See*, *e.g.*, 2d Bradner Decl. ¶¶ 114-29 (JA 7:3919-26; 3d Schulzrinne Decl. ¶ 3 (JA 7:4021).

We conclude that the government's "traffic mirroring hypothetical" is in direct conflict with the technical and economic realities of running large international computer networks. In order to implement either mirror ports or an optical splitter, the government would have to work closely with ISPs. We conclude no network operator would choose configuring mirror ports and filters on Core Routers over using an optical splitter - assuming it is even possible to implement mirror ports in the locations required.

The government tries to diminish the likelihood of the optical splitter copy-all-then scan approach. The government's expert suggests that there are "practical considerations" that weigh against it. 2d Schulzrinne Decl. ¶ 26 (JA 6:3419) ("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"). But Dr. Schulzrinne is mistaken. Using a mirror port rather

than an optical splitter introduces far more significant practical risks.

Mirror ports are not "operationally speaking ... imperceptible to the carrier." 2d Schulzrinne Decl. ¶ 24 (JA 6:3418). Mirror ports require additional physical ports. Adding ports requires real capital expenditure outlays, and also consumes chassis slots, which are a finite resource in the hardware of the physical routers, or "router chassis." No amount of money can add more slots to a router chassis. Mirror ports also consume significant processing power on Core Routers, cost employee time to manage, require changes in processes and procedures, and introduce a significant chance of interrupting customer traffic — literally putting the ISP's business at risk.

In contrast, an optical splitter is extremely reliable as it consumes no power, has no software, and cannot slow traffic. The risk and cost of an optical splitter is a fraction of the risk and cost incurred in maintaining mirror ports with filters. In addition, optical splitters do not require significant disclosures to network operator employees about what communications are being surveilled. By contrast, to implement port mirroring in a typical tier one ISP, several dozen to well over 100 engineers, technicians, and even outside vendors, would be able to see the list of selectors configured into a router's whitelists and blacklists. It is likely that many of the ISP's staff are neither US citizens nor residents.

Thus, given the necessity of serving customers reliably and cost effectively to ensure a viable business, we conclude no network operator would choose configuring mirror ports and filters on Core Routers over using an optical splitter. The government's arguments to the contrary simply misunderstand the technical reality of traffic mirroring.

B. THERE ARE NUMEROUS TECHNICAL IMPEDIMENTS TO USING TRAFFIC MIRRORING TO AVOID WIKIMEDIA'S COMMUNICATIONS

Even if we assume that the government (and the ISPs it works with) foregoes the use of optical splitters and chooses to use traffic mirroring as Dr. Schulzrinne hypothesizes, traffic mirroring on a Core Router could not be implemented as the government describes for several reasons: (1) Core Routers were not designed to reliably apply large scale traffic mirroring; (2) "traffic encapsulation" would make the government's proposed technique impossible in most circumstances; and (3) implementing blacklists and whitelists of large volume, low-interest websites is not viable due to the widespread use of Content Distribution Networks (CDNs).

Before we address each of these impediments, a brief explanation of Access Control Lists (ACLs) is warranted. The district court focuses on whitelists and blacklists as the means to avoid all of Wikimedia's communications, which would be implemented through ACLs on Core Routers. ACLs "perform packet filtering to

control which packets move through a network and to where." Security Configuration Guide: Access Control Lists, Cisco IOS XE Release 3S, CISCO, https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_data_acl/configuration/xe-3s/sec-data-acl-xe-3s-book/sec-access-list-ov.html. ACLs are "counted" by the number of lines in the configuration. Generally speaking, a longer ACL (more lines) uses more of the Core Router's processing power than a shorter ACL (fewer lines). But Core Routers face numerous technical limitations in implementing ACLs. It is highly unlikely, if not virtually impossible, that the government could implement combinations of whitelists and blacklists in the manner described by the government to avoid copying, collecting, or otherwise reviewing any of Wikimedia's communications.

In the government's hypothetical technique, ACLs are configured on mirror ports to limit the packets being forwarded to the government's surveillance device. Under this technique, the Core Router examines fields in the header of each packet being sent to the mirror port, then takes different actions based on whether the specified packet header fields match the ACL or not.

The district court recognized three fields in the packet header used to filter in the government's hypothetical technique: (1) protocol number, (2) port number, and (3) IP address. SJ Op. at 11 (JA 7:4083). It is important to note that while ACLs on

Core Routers can filter on more than the three fields listed, there are limitations. For example, it is not possible for a Core Router to filter on email addresses in the payload of a packet.

The district court examined two possible actions for the ACLs to take on mirror ports: (1) a "blacklist" which denies all packets that match the filtered field from being passed to the surveilling entity; and (2) "whitelists" which allows all packets that match the filtered field to be passed to the surveilling entity. SJ Op. at 13 (JA 7:4085). It is possible to use both blacklists and whitelists at the same time. For example, an ACL can be configured to whitelist (allow) a large range of IP addresses, then also blacklist (deny) any packets with certain port numbers. The use of ACL whitelists and blacklists is central to the government's hypothetical technique that the Upstream program could readily be implemented in a way to avoid all of Wikimedia's communications.

But the technical realities of implementing the government's theory leads us to conclude, like Mr. Bradner did, that the government's hypothetical is not implementable in the real world. *See* 2d Bradner Decl. ¶¶ 6, 55-148, 154-55 (JA 7: 3884, 3899-3939).

1. THE DISTRICT COURT ERRED IN ACCEPTING THE TECHNICAL FEASIBILITY OF USING ACCESS CONTROL LISTS TO AVOID ALL OF WIKIMEDIA'S COMMUNICATIONS

Configuring ACLs on Core Routers as suggested in Dr. Schulzrinne's hypothetical is technologically infeasible. The district court says the government's hypotheticals show that "there is a technological method by which the NSA could conduct Upstream surveillance on a circuit transporting International internet communications without copying, collecting, or otherwise reviewing any of Wikimedia's communications that traverse that path." SJ Op. at 17 (JA 7:4089). Dr. Schulzrinne goes into significant detail of how he envisions such a scheme would look, saying:

> "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)."

2d Schulzrinne Decl. ¶ 35 (JA 6:3423). But this academic thought experiment could not be feasibly implemented on actual networks. While it is possible to find a list of IP addresses for webmail providers, the list would be too long to configure on nearly any ISP's Core Routers. Using Gmail as an example, when a user types

"mail.google.com" into their browser to read email, nearly 2,000 IP addresses can be returned for that one hostname on that one provider.

Thus this argument runs into the same problem that plagues each of the government's responses to Mr. Bradner's conclusions. Dr. Schulzrinne's hypothetical technique makes assumptions which may seem reasonable based on equipment feature lists found online, but are simply not possible to execute in real world conditions. Core Routers were not designed for the kind of extensive granular filtering that would be necessary to implement ACLs in the manner suggested by Dr. Schulzrinne. Implementing ACLs that would avoid all of Wikimedia's communications while still collecting targeted selectors in the manner described in public disclosures would pose a high risk to the operation of those networks.

The hypothetical technique the government proposes for avoiding all of Wikimedia's communications is unconvincing. We conclude that it is highly unlikely, if not virtually impossible, that the government implements combinations of whitelists and blacklists in the manner described above to avoid copying, collecting, or otherwise reviewing any of Wikimedia's communications.

2. THE WIDESPREAD USE OF TRAFFIC ENCAPSULATION TECHNOLOGIES WOULD PREVENT THE GOVERNMENT FROM FILTERING TRAFFIC USING ACLS

Another problem with the hypothetical technique credited by the district court

is that nearly all tier one and similarly sized ISPs use encapsulation technologies, the most common of which is called Multiprotocol Label Switching (MPLS). These technologies essentially place a wrapper around data packets. The Core Router never "sees" the packet header, and therefore ACLs cannot filter based on any of the fields discussed by the district court. As the majority of tier one ISPs encapsulate backbone traffic with MPLS (or similar), it is not possible to use the government's hypothetical technique to capture target traffic while filtering out Wikimedia's communications.

Therefore, the district court's opinion concerning the technical feasibility of the government's argument is based on false assumptions.

3. CONTENT DISTRIBUTION NETWORKS MAKE IMPLEMENTING ACLS ON CORE ROUTERS TO AVOID LOW INTEREST WEBSITES TECHNOLOGICALLY INFEASIBLE

Lastly, the widespread use of CDNs make the government's proposed hypothetical to avoid high-traffic, but low-interest websites technologically infeasible. The district court credits Dr. Schulzrinne when he suggests filtering "traffic from high-volume websites such as Amazon.com and Wikipedia" could be used to avoid Wikimedia's communications. 2d Schulzrinne Decl. ¶ 22 (JA 6:3418). The district court's opinion finds that under the hypothetical technique, there is a "technological method" that the government could avoid all of Wikimedia's communications by using "a combination of whitelisting and blacklisting" to

exclude low interest communications, which might include Wikimedia's. SJ Op. at 17 (JA 7:4089).

Dr. Schulzrinne suggests using Alexa (https://www.alexa.com/topsites/) as a source for tracking the most popular websites that may be of low interest. 2d Schulzrinne Decl. ¶ 41 (JA 6:3426). He states "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." *Id.* To perform these actions, Dr. Schulzrinne assumes "the IP addresses of the servers that host Amazon.com, or Wikipedia.org, must remain unchanging if online shoppers, or Wikipedia's readers and contributors, are to reach them over the Internet."[5] 1st Schulzrinne Decl. ¶ 33 (JA 1:0734).

This assumption is incorrect. In reality, the IP address of most content on the Internet today, including websites, webmail, social media, and chat platforms, is neither fixed nor exclusive to that content. This is because most large websites, including 48 of the 50 websites listed on the Alexa US top 50 are served from Content Distribution Networks (CDNs). *See Top Sites in United States*, Alexa, (Jul. 2, 2020) https://www.alexa.com/topsites/countries/US.

[5] While www.wikipedia.com is not served on a commercial CDN, the district court's conclusion is based on the notion there are many high traffic, low interest websites which have very few, unchanging IP addresses. This conclusion is mistaken.

A CDN is a collection of servers distributed over a large area, typically globally. Four of the largest CDNs by traffic (Akamai, Google, Netflix, Facebook) each have tens or hundreds of thousands of servers, deployed in over 100 countries. Using a CDN allows a content owner to benefit from the performance, scale, and reliability of a global server network.

Generally, a Tokyo-based user would be directed to a server in Tokyo, while a Miami-based user would be directed to a server in Miami. For example, the federal courts website administrator engaged the CDN provider Akamai to serve www.uscourts.gov. When an end user goes to www.uscourts.gov, Akamai chooses the optimal server to serve the web page based on network and server conditions.

Because CDNs dynamically supply content based on the location of the end user, network conditions, server load, and many other variables, the CDN will constantly change the IP address it uses to serve each website. These changes are frequent — Akamai, for example, guarantees a website's IP address will remain stable for only 20 seconds. Additionally, CDNs often serve many websites from a single IP address. *See* John Dilley, Bruce Maggs, et. al., Global Deployment of Data Centers, IEEE Internet Computing (Sept./Oct. 2002) https://people.cs.umass.edu/~ramesh/Site/PUBLICATIONS_files/DMPPSW02.pdf ; Akamai, *What Does CDN Stand For? CDN Definition*,

https://www.akamai.com/us/en/cdn/what-is-a-cdn.jsp (last visited June 18, 2020).

As a result of CDNs, Dr. Schulzrinne's hypothetical technique is impossible for at least three reasons:

First, the number of lines an ACL would have to contain to either whitelist or blacklist even one CDN, let alone the more than a dozen CDNs used by the websites in the Alexa US top 50, alone, is far more than Core Routers on the Internet can handle. Second, it is not possible to update the router configurations as quickly as CDNs add or change IP addresses. Third, ACLs cannot be used to allow or deny only specific websites, social media platforms, chat rooms, etc. if those sites use any of the CDNs which serve multiple websites from a single IP address.

Therefore, the district court's opinion concerning the technical feasibility of the government's argument is based on false assumptions.

III. CONCLUSION

Wikimedia is challenging the constitutionality of the alleged interception of at least some of its communications. *Amici* express no opinion on that underlying legal question. But we agree with Wikimedia's expert, that as a technical matter, it is "virtually certain" the government is not avoiding the copying or interception of all of Wikimedia's communications under the Upstream surveillance program. *See, e.g.*, 2d Bradner Decl. ¶¶ 114-15 (JA 7:3919-20). For the reasons discussed above,

amici submit that this Court should reverse the district court ruling dismissing the case.

DATED: July 8, 2020

Respectfully submitted,

/s/ Jonathan H. Blavin
JONATHAN H. BLAVIN

Jonathan Blavin
Munger, Tolles & Olson LLP
560 Mission St., 27th Floor
San Francisco, California 94105
(415) 644-6911

Counsel for *Amici Curiae*

APPENDIX A
List of *Amici Curiae*

All *amici* sign in their personal capacity *only*, and the following titles and employer affiliations are provided for identification purposes only. No *amici* are acting as representatives of their employers.

- Patrick Gilmore, Founder, Deep Edge Technologies
- Nat Meysenburg, Technologist, Open Technology Institute
- Bruce Schneier, Fellow, Harvard Kennedy School and Chief of Security Architecture at Inrupt, Inc.
- Chad W. Milam, Network Architect and Industry Executive with 26 years of experience
- John Kristoff, Network Architect, DePaul University
- Michael Young, CISM CISA, Board Member, New York Information Systems Security Association, US and International Board Adviser, Information Systems Security Association

- Susan Forney, Network Engineer, Hurricane Electric
- J. Alex Halderman, Professor of Computer Science and Engineering and Director, Center for Computer Security and Society, University of Michigan
- Stephen Wilcox, President & Founder, IX Reach
- David Crocker, Brandenburg InternetWorking
- Joel Jaeggli, Fastly
- Stephen Farrell, Trinity College, Dublin
- Roger Dingledine, The Tor Project
- Wesley George, Vice President of Networking and CTO, DataBridge Sites, LLC.
- Richard A Steenbergen, CEO, Petabit Scale
- Allison Nixon, Chief Research Officer, Unit 221B
- Mark Rasch, Professorial Lecturer of Law, GW Law School,
- Christian Kaufmann, Network Architect with 20 years of experience
- Huanhuan Jezzibell Gilmore, Chief Commercial Officer, PacketFabric, Inc.
- Kristin Berdan, Research Fellow, Center for Long-Term Cybersecurity, University of California Berkeley
- L. Sean Kennedy, Itaunas Telecom Consulting
- Joseph Lorenzo Hall, Senior Vice President, Strong Internet

CERTIFICATE OF COMPLIANCE WITH RULE 32(a)

IT IS HEREBY CERTIFIED:

1. That the foregoing Brief *Amicus Curiae* complies with the word count limitation set forth by Rule 29(d) because this brief contains 4,423 words, excluding the parts of the brief exempted by Rule 32(f).

2. This brief complies with the typeface requirements of Fed. R. App. P. 32(a)(5) and the type style requirements of Fed. R. App. P. 32(a)(6) because this brief has been prepared in a proportionally spaced typeface using Microsoft Word 2016 in 14-point Times New Roman.

DATED: July 8, 2020 MUNGER, TOLLES & OLSON LLP

By: */s/Jonathan H. Blavin*
Jonathan H. Blavin

Attorney for Amici Curiae

45151633.1

CERTIFICATE OF SERVICE

I hereby certify that I electronically filed the foregoing with the Clerk of the Court for the United State Court of Appeals for the Fourth Circuit by using the appellate CM/ECF system on July 8, 2020.

I certify that all participants in the case are registered CM/ECF users and that service will be accomplished by the appellate CM/ECF System.

DATED: July 8, 2020 MUNGER, TOLLES & OLSON LLP

By: */s/Jonathan H. Blavin*
Jonathan H. Blavin

Attorney for Amici Curiae

45151633.1