Loopholes for Circumventing the Constitution: Warrantless Bulk Surveillance on Americans by Collecting Network Traffic Abroad

Working Paper. Last updated June 27, 2014.

Axel Arnbak[1] and Sharon Goldberg[2]

[1] Berkman Center for Internet & Society, Harvard University, Cambridge, MA 02138.
[2] Computer Science Department, Boston University, Boston, MA 02215.

Abstract. In this multi-disciplinary paper, we reveal interdependent legal and technical loopholes that intelligence agencies of the U.S. government could use to circumvent constitutional and statutory safeguards for U.S. persons. We outline known and new circumvention techniques that can leave the Internet traffic of Americans as vulnerable to surveillance, and as unprotected by U.S. law, as the Internet traffic of foreigners.

Keywords: surveillance, privacy, FISA, Executive Order 12333, network protocols, DNS attacks, BGP attacks.

1 Introduction

As the general public and the media becomes overloaded by the string of recent N.S.A. revelations, the academic community has the important task of precisely describing the legal and technical realities under which these programs operate, and to offer informed recommendations on how to overcome the current status quo of apparent surveillance overreach. In this multi-disciplinary paper, we reveal interdependent legal and technical loopholes that intelligence agencies of the U.S. government could use to circumvent 4th Amendment and statutory safeguards for Americans. We focus on the legal and technical dimension of network surveillance by intelligence agencies in the data collection phase, rather than during data retention, or further analysis once data is collected. Our central hypothesis is that there are several loopholes that these authorities can exploit to conduct largely unrestrained surveillance on Americans by collecting their network traffic *abroad*.

Legal Loopholes. In Section 2 we start by describing the current U.S. regulatory framework for intelligence gathering. From public and until-recently secret primary legal sources, three regimes can be distinguished, based on *where* the surveillance is conducted, and *who* it targets:

1. Surveillance of domestic communications conducted on U.S. soil under s.215 of the "Patriot Act";
2. Surveillance of foreign communications conducted on U.S. soil under the "Foreign Intelligence Surveillance Act"; and

3. Surveillance conducted entirely abroad under "Executive Order 12333" (EO 12333) and its minimization policies, notably U.S. Signals Intelligence Directive 18 ("USSID 18"). USSID 18 was drafted and approved within the Executive branch with minimal Congressional or Judicial oversight.

The first two regimes are overseen by all three branches of the U.S. government, and currently under scrutiny by the government, media and the general public. The third regime, however, is solely the domain of the Executive branch and has largely been ignored by the public and other branches of Government in recent months, especially since relevant legal documents related to EO 12333 remain classified or redacted. However, according to the N.S.A., this third regime under EO 12333 is the 'primary legal authority' for its operations [5, p. 2-3]. Thus, it deserves more attention and careful scrutiny.

Working with primary legal sources, many of which have only recently been made public and are still redacted on key issues, we make the following central observation. A surveillance operation falls within the EO 12333 regime when it *presumes* two connected criteria: it does not *intentionally target a U.S. person*, and is *conducted abroad*. If an intelligence agency can construct plausible presumptions that these two criteria have been meet, then the permissive legal regime under EO 12333 can be applied to the surveillance operation. The surveillance is then considered to affect non-U.S. persons, and 4th Amendment protections can thus be circumvented even if the operation primarily affects Americans. Our main hypothesis is therefore that *there is loophole for surveillance on Americans from abroad* resulting from the following interdependence: (1) the complete absence of legal protection for non-U.S. persons under the U.S. regulatory framework [32,33] creates 'foreignness'-presumptions under EO 12333 and (2) the technical realities of modern Internet communications.

Technical Loopholes. At first blush, one might suppose that a surveillance operation conducted *abroad* should have no impact on the privacy of Americans. However, in Section 3 we discuss why the technical realities of the Internet mean that American's network traffic can easily be routed or stored abroad, where it can then be collected under the permissive legal regime of EO 12333. Indeed, we already know of surveillance programs that have exploited this legal loophole. The revealed MUSCULAR/TURMOIL program, for example, illustrates how the N.S.A. presumed authority under EO 12333 to acquire traffic between Google and Yahoo! servers located on foreign territory; this program allegedly collected up to 180 million user records per month abroad, including those of Americans [17].

We also discuss other technical means an intelligence agency can exploit the legal loopholes under EO 12333. Instead of eavesdropping on *intradomain* traffic (*i.e.,* data sent within a network belonging to a single organization, as in the MUSCULAR/TURMOIL program), these loopholes can be exploited in the *interdomain* setting, where traffic traverses networks belonging to different organizations. We explain why interdomain routing with BGP can naturally cause traffic originating in a U.S. network to be routed abroad, even when it is destined for an endpoint located on U.S. soil. We also discuss why core Internet

protocols – BGP and DNS – can be *deliberately* manipulated to force traffic originating in American networks to be routed abroad. We discuss why these deliberate manipulations fall within the permissive EO 12333 regime, and how they can be used to collect, in bulk, all Internet traffic (including metadata and content) sent between a pair of networks; even if both networks are located on U.S. soil (*e.g.,* from Harvard University to Boston University).

We do not intend to speculate on whether or not the intelligence community is exploiting the interdependent technical and legal loopholes that we describe in this paper. Instead, our aim is to broaden our understanding of the possibilities at hand. Our analysis suggests that, without a fundamental reconsideration of the lack of privacy and due process safeguards for foreigners, current surveillance legislation opens the door for unrestrained surveillance on Americans from abroad.

Paper Organization. To examine our central hypothesis, we combine descriptive, internal legal analysis with threat modeling from computer science. This paper presents the status of our work in progress. Section 2 identifies legal loopholes in the three legal regimes that form the regulatory framework for network surveillance for intelligence agencies. Section 3 discusses the technical details of how network protocols can be exploited to circumvent the legal protections for Americans. We conclude with a brief description of the policy recommendations that we are developing as part of future work, and a reflection on the fundamental problems of legal protections that follow national borders in a global communications network. Our research method also offers new insights for normative policy evaluation and a multi-disciplinary framework for further research.

2 Loopholes in the Legal Framework

A recurring conundrum for regulation of global communications networks is that the application of law is, ultimately, tied to jurisdiction. For centuries, jurisdiction has been determined primarily by physical and geopolitical borders, or the space that states consider sovereign territory. Because networked communication does not necessarily respect physical and geopolitical borders, transnational surveillance (*i.e.,* surveillance conducted from one country, directed towards users in another country) presents us with one of the most urgent examples of this conundrum [32].

In this section, we use recently revealed and declassified primary legal sources to describe and contextualize the U.S. legal framework for network surveillance by intelligence agencies. We will demonstrate that three legal regimes outlined in Section 1 can be distinguished based on two main criteria: *who* is targeted (US person or not), and *where* the communication is taking place (on U.S. territory, or abroad). Our focus is on the poorly-understood third regime under EO 12333, which authorizes the N.S.A. to conduct largely unrestrained surveillance operations on foreign soil. Since the third regime covers operations on foreign soil that are not covered by the first two legal regimes (s.215 of the "Patriot

Act" and the "Foreign Intelligence Surveillance Act (FISA)" [30, p.3]), we must start by analyzing the types of operations that fall under those two legal regimes. Our discussion will also highlight the differences in legal protection under each regime, and why the outlook on reform fundamentally differs. We then move on to discussing the third legal regime (EO 12333) in detail, and find that it applies when surveillance does not 'intentionally target a U.S. person' and is conducted abroad, regardless of whether or not the operation affects millions of communications records of Americans.

2.1 First Regulatory Regime: Domestic Communications, Surveillance Conducted on US Soil.

Some intelligence surveillance operations target domestic communications on U.S. soil. The legal framework of this class of operations is relatively well-known. Under section 215 of the "Patriot Act", intelligence agencies can request a warrant at the FISA Court for 'tangible things' that are 'relevant' to authorized terrorism or counterintelligence investigations. This s. 215 was adopted in its current form soon after the attacks of 9/11, and significantly broadened the legal authority for the N.S.A. to conduct domestic surveillance.

Meanwhile, a well-covered program operating under this legal authority is the bulk collections of Americans' telephone records under the so-called Verizon Metadata Program. Immediately after the 9/11 attacks, U.S. President Bush arranged for the voluntary provisions of these records by all the major U.S. telecommunications providers. Upon a 2005 disclosure in the press of the program, one company asked the government to obtain a warrant from the FISA Court. Since 2006, the Court has granted the warrants on a rolling basis, including so-called 'gag' orders that prevent the companies from disclosing the bulk metadata requests to customers or the wider public [23].

With the details of the telephony metadata programs revealed after nearly twelve years, scholars have argued that the program violates both the provisions of the Patriot Act and the Constitution [12]. In U.S. Congress, proposals have been initiated in the U.S. Congress to address the surveillance overreach [29]. Furthermore, several court cases are pending in different judiciary circuits, implying that the U.S. Supreme Court will finally rule on the matter in the not so distant future. Regardless of the short-term outcome of these legal and political debates, on the long term three branches of government (*i.e.,* the Executive, the Legislative and the Judiciary) will be involved in establishing checks and balances between the three of them, as well as legal protections for Americans against surveillance overreach by the Executive branch.

2.2 Second Regulatory Regime: Foreign Communications, Surveillance Conducted on U.S. Soil.

The second regulatory regime covers a class of surveillance operations conducted on U.S. soil, regulated by the 1978 Foreign Intelligence Surveillance Act ('FISA'). These operations are aimed at obtaining 'foreign intelligence information', a

broad term that includes information 'relating to the foreign affairs of the U.S.' (cf. art. 1801(e)(2) of FISA). This includes economic and political surveillance of foreign governments, corporations, media organizations and citizens [33].

First we provide an overview of this second regime. We then describe which surveillance operations (particularly those with an international aspect) are covered under FISA. We find that other international surveillance operations, including those enabled by the network protocol manipulations we present in Section 3, fall under the more permissive third legal regime for surveillance discussed in Section 2.3. Finally, we discuss the legal protections afforded to Americans under FISA as well as FISA reform.

2.2.1 Overview of the Second Regulatory Regime under FISA.

FISA and the FISA Court were introduced in 1978 by U.S. Congress, in response to domestic surveillance overreach and the reform proposals by the Church Committee [33]. In 2008, FISA was amended and significantly broadened by U.S. Congress with the FISA Amendments Act ('FAA'). The FAA introduced section 702, which allows for *warrantless* surveillance of foreign communications conducted on U.S. soil, as long as these operations do not 'intentionally target U.S. persons'. That is, s. 702 does not require warrants to be issued for a specific case based on a *particularized* probable cause. Instead, the FISA Court approves of *generalized* 'targeting' and 'minimization' procedures on any data that is collected; these procedures are intended to mediate U.S. person privacy concerns, and have remained classified until recently [32].

For years, FISA and especially its s. 702 have been criticized for providing legal loopholes for warrantless political and economic surveillance on U.S. lawyers, NGOs, journalists and corporations communicating internationally through U.S. Internet companies [33]; the media reports in December 2005, around warrantless wiretapping in bulk from the Internet backbone at an AT&T switch [28], have highlighted some of this tension. Nonetheless, U.S. Congress passed FAA after the AT&T revelations and extended the validity of the FAA for another five years on 31 December 2012, one day before the sunset deadline. Two months later, on 26 February 2013 in the case 'Clapper v. Amnesty International', the U.S. Supreme Court denied several U.S. organizations a right to claim that the privacy of their international communications was violated by s. 702 *on procedural grounds*. In what appeared to be the final ruling on the constitutionality of s. 702 for the foreseeable future, a 5-4 majority argued that these organizations were merely 'speculating', and could not prove that their communications had actually been intercepted [6]. Justice Breyer, on behalf of the minority, noted in his dissent that s. 702 prohibits the same applicants to actually gain knowledge of the surveillance itself because of national security secrecy, and that the broad authorities probably existed for a reason.

The political debate and the issue of legal standing have shifted considerably since June 2013, when it became clear that s. 702 indeed serves as the legal basis for many operations, among them 'UPSTREAM' and 'PRISM' [13]. Moreover, several of the classified targeting and minimization procedures under s. 702 have

been leaked or declassified [2, 3]. Both revelations have spurred the N.S.A. to confirm that a principle use of s. 702 is compelling assistance from U.S. Internet companies for warrantless surveillance [5, p. 4].

This new dynamic enables a unique insight into classified and generous interpretations of the legal provisions in FISA made by the intelligence community and the FISA Court [13]. Before we dive into the details of FISA, we mention that FISA also contains s. 703 and s. 704, that regulate surveillance intentionally targeting U.S. persons located abroad. These sections are outside the scope of this paper, since our focus is on surveillance operations on Americans located in the U.S., with surveillance conducted on foreign soil. As an aside, Donohue has observed that the warrant requirements in these sections have been circumvented by applying s. 702 criteria to the collection phase, and then seeing whether collected data is of use for further processing after the fact [13, p.26].

2.2.2 Scope of the Second Regulatory Regime under FISA: The 1978 'Electronic Surveillance' Definition

All communications surveillance operations that constitute 'electronic surveillance', as defined s. 1801(f) of FISA, fall within the scope of FISA (cf. 18 U.S.C. s.2511(2)(f); 50 U.S.C. s.1812(a)). The definition has largely remained intact since 1978. To acquire the content of 'wired communications', surveillance only falls within the FISA definition when authorities 'intentionally target a U.S. person' (s. 1801(f)(1)), or when the acquisition is conducted on U.S. soil (s. 1801(f)(2)). Importantly, when authorities conduct targeted surveillance from abroad, even if they know that both 'sender and all intended recipients are located in the U.S.', then only 'radio' (*i.e.,* wireless) communications fall within the FISA definition of 'electronic surveillance' (s. 1801(f)(3)). The FISA definition only mentions communications 'content', but not 'metadata' (location, time, duration, identity of communicants, etc.), which in itself gives rise to privacy concerns that we will not further discuss here. Relevant for our purposes, is the observation that operations on 'wired communications', when conducted abroad, only fall within the scope of FISA if they 'intentionally target a U.S. person'.

Intentionally Targeting U.S. Persons. 'Intentionally targeting a U.S. person' constitutes 'electronic surveillance' under FISA (s. 1801(f)(1)). However, 'intention' and 'targeting' are not defined in FISA, leaving the concepts open to generous interpretation by authorities in classified 'targeting' and 'minimization' procedures. Apart from providing clarity that bulk surveillance is not regarded as intentional targeting (we discuss this further when we look at legal protections from U.S. persons under FISA), the disclosure of these procedures has revealed two important new facts related to surveillance operations conducted abroad. Firstly, conducting the surveillance abroad creates the presumption that the surveillance targets a non-U.S. person [2, p. 3-4]. Secondly, the 'targeting procedures' do not provide any due diligence requirement or duty of care to establish the identity of parties on either side of a communication [2, p.3-4] [3]. This implies that unless a communicant is known to be a U.S. person, the procedures

consider the communicant to be a non-U.S. person. In other words, authorities have a strong incentive to conduct surveillance abroad: legal protections offered to U.S. persons under FISA can be circumvented, and a more generous legal regime applies to the data collection itself.

Installing a Device. Of particular interest to our analysis is preparing a communications infrastructure for surveillance, *e.g.,* via network protocol manipulations that modify the flow of network traffic as described in Section 3.2. FISA has a clause on 'installing a device for that purpose in the United States', which can be understood as making a communications infrastructure 'ready' for surveillance, but it only covers *'other than wire or radio communication'* (s. 1801(f)(4)). The U.S. Congressional Research Service gives 'a hidden microphone' as an example of such 'other communication' [30, p.7]). Even if advanced protocol manipulations or contemporary active attacks such as injecting malware or installing backdoors to enable surveillance would fall under the 1978 'Installing a device' definition, when it concerns 'wired communications' such operations apparently falls outside FISA altogether, both when conducted on U.S. soil or abroad.

In addition to such close textual analysis, another legal analytical tool supports this observation; the current consensus amongst scholars and policymakers is that the intention with the 'electronic surveillance' definition of Congress, both in 1978 and today, has always been to carefully 'exclude a lot of what the N.S.A. is actually doing' [14]. The definition has not been updated much since 1978, which to further supports the idea that the legislator did not intend to update FISA to provide further legal protections to U.S. persons in modern communications environments. Thus, both textualist and originalist legal interpretative doctrines support our observation, and we agree with [13, p.26] that the case for a dynamic interpretation of the 1978 definition is weak.

In short, apart from passive surveillance, even active attacks such as the advanced network protocol manipulations we describe in Section 3 only fall under the 'electronic surveillance' definition in FISA when U.S. persons are intentionally targeted. Our observations are supported by recent revelations on malware operations in Der Spiegel (further discussed in the next section). [4]

2.2.3 Legal Protections for U.S. Persons under FISA.

Applicability of FISA to a surveillance operation is relevant for Americans, because the statue contains some important legal protections for U.S. persons intentionally targeted. For instance, the statute explicitly states that the 4th Amendment applies to surveillance operations under FISA (cf. s.1881(b)(5)) and a narrow set of four surveillance operations is explicitly prohibited. As discussed, surveillance under s. 702 may not intentionally target a U.S. person; for those operations s. 703 exists. Another example is the **'reverse-targeting'** prohibition of s.1881(b)(2), which holds that authorities may not intentionally target a non-U.S. person under a s. 702 if the actual purpose of the operation is to target a U.S. person. By contrast, the third legal regime under EO 12333 explicitly al-

lows for intentional targeting of U.S. persons, when certain conditions discussed in the next section are met.

Nonetheless, serious loopholes exist for surveillance conducted within the bounds of FISA. One of the most-discussed loopholes is when U.S. persons have not been 'intentionally targeted' but instead affected by a surveillance operation, *e.g.,* a bulk intercepts on the Internet backbone on U.S. soil under the 'UPSTREAM' program. Instead of promptly destroying such data, generous exemptions exist to nonetheless use the 'incidentally' or 'inadvertantly' collected information of the affected U.S. person, including when a 'foreign intelligence' interest is created in the data sometime after its collection, or when the information could be relevant for cybersecurity (incl. cyber-offense) purposes [3].

More generally, the targeting and minimization procedures seem to have introduced a new category of surveillance specifically aimed acquiring information *about persons*. (For example, two communicants that chat *about* a subject, like Angela Merkel, which is part of an N.S.A. 'selector'.) Such surveillance is not considered to intentionally target specific communicating parties, and hardly enjoys protection even if it affects U.S persons. The information collected through such operations may be further analyzed and disseminated to other agencies as long as the identity of U.S. persons implicated are redacted in a way 'that the information cannot be reasonably connected with an identifiable U.S. person' [3, s.6]).[3] A more complete analysis of the targeting and minimization procedures can be found in [13], along with a critical assessment of the role of the FISA Court.

2.2.4 Reforming the Second Regulatory Regime.

FISA and FAA have serious implications for the privacy rights of Americans. And current reform proposals, including the proposed USA Freedom Act, pay far too little attention to the loopholes in the antiquated 1978 FISA definition of 'electronic surveillance' and the permissive workarounds for the restrictions on 'intentionally targeting U.S. persons'. Nonetheless, adopting a long term perspective on reform, the FISA and FAA statutes have been approved by the U.S. Congress, while the targeting and minimization procedures have been approved by the FISA Court. In response to the recent disclosures, proposals have been made to reform this legal regime, including tightening the s. 702 loopholes and making hearings before the FISA Court adversarial by allowing a 'civil liberties advocate' to defend privacy interests. As with domestic surveillance, all three branches of government will be involved in long term FISA reform. As such, the barriers to strengthening privacy rights of Americans are mostly political, not institutional. We will see that this is not the case in the third legal regime.

[3] Redacting information as a means of protecting identities comes with its own privacy issues; see [25] for more discussion.

2.3 Third Regulatory Regime: Surveillance Conducted on Foreign Soil.

Electronic surveillance conducted abroad is by and large regulated by "Executive Order 12333" (EO 12333). Surveillance policies regulated under this regime are designed and adopted solely within the Executive branch. The N.S.A. recently acknowledged that EO 12333 is "the foundational authority by which N.S.A. collects, retains, analyzes, and disseminates foreign signals intelligence information" [5, p. 2-3]. Even so, EO 12333 and its underlying policies are hardly discussed in policy and scholarly circles. This may be explained by the conventional secrecy surrounding national security policy: understanding how exactly EO 12333 surveillance is shaped was practically impossible outside the intelligence community, until recently.

As with FISA, the Snowden disclosures and subsequent developments provide a unique opportunity to gain understanding of EO12333 and, in particular, the policies that specify the authorities it provides. EO 12333 itself is very general, as is the 2007 Department of Defence Directive 5240.01 that lays a general framework of rules for intelligence conduct based upon the order. DoD Directive 5240.1-R of 1982, adopted over three decades ago, contains further principles on 'DoD activities that may affect U.S. persons'. Aforementioned documents are not particularly interesting in themselves, but form the basis of U.S. Signals Intelligence Directive 18 ("USSID 18"), that is lower in legal hierarchy but becomes fairly specific on actual surveillance principles.

Until recently only a 1993 version of USSID 18 was de-classified. Most probably in response to the MUSCULAR revelations on 30 October 2013 [17], the authorities released a 2011 version of USSID 18 on 18 November 2013 [7] that remains redacted on critical parts, as we will see. Noting that the 1982 DoD Directive, until this day, contains a completely classified Annex A particularized for N.S.A. conduct that may affect U.S. persons, we focus our analysis on this recently declassified 2011 version of USSID 18. With regard to actual operations, the public has learned how the N.S.A. assumed authority under EO 12333 to acquire communications within Google and Yahoo! networks because the operation was conducted on foreign territory [17], collecting up to 180 million user records per month, regardless of nationality (we discuss this program is Section 3.1.)

Nevertheless, unlike under FISA, critical parts of the underlying policies are still classified, or heavily redacted. In this section, we will analyze what is publicly known about this third regime and indicate the remaining knowledge gaps. We first discuss the scope of EO 12333, and when it applies to advanced network surveillance methods. We then describe how U.S. intelligence authorities enjoy broad and largely unchecked legal authority when conducting surveillance abroad, and how legal protection offered to Americans under EO12333 are substantially lower than under the other regimes. Finally, we point at fundamental institutional barriers in the U.S. Constitution to long term reform of EO12333 policies, regardless of their serious impact on Americans' privacy.

2.3.1 Scope of the Third Regulatory Regime under EO 12333: Electronic Surveillance Conducted Abroad.

As discussed in the Section 2.2, electronic surveillance falls within the EO 12333 regime when it is conducted on foreign soil, and when it does not fall within the 1978 FISA definition of 'electronic surveillance'. Or as the N.S.A. recently put it, when surveillance is "conducted through various means around the globe, largely from outside the United States, which is not otherwise regulated by FISA." [5, p. 2-3]. [4]

While FISA surveillance is conducted from U.S. soil, EO 12333 surveillance is mostly conducted abroad. EO 12333 presumes that network traffic intercepted on foreign soil belongs to non-U.S. persons (cf. s. 9.8 & 9.18.e.2 of USSID 18 defining 'foreign communications' and 'U.S. person'). Companies and associations are also considered in the EO 12333 definition of U.S. persons. These entities may be assumed to be non-U.S. persons if they have their headquarters outside the U.S. Even when it is known to the N.S.A. that a company is legally controlled by a U.S. company, it may be assumed a non-U.S. person. Taken together, the rules for presuming a non-U.S. person under this regime are permissive on the individual-, group- and organizational levels.

Installing a Device. We return to the question of 'installing a device', to understand how EO 12333 regulates the network protocol manipulations we describe in Section 3.2. These manipulations fall under EO 12333. However, on top of the 1978 FISA definition of 'electronic surveillance', neither EO 12333 nor the 2011 update of USSID 18 further specify what 'installing a device' means today. It is not covered in the definitions of 'collection' (section 9.2 USSID 18), 'interception' (section 9.11 USSID 18) nor in the definition of 'electronic surveillance' (section 9.7 USSID 18). The definition of 'installing a device' to enable surveillance could possibly be redacted in USSID 18 or further specified in a still-classified guideline. A post-Snowden N.S.A. memorandum does not provide any clarity [5, p. 2-3]. To the contrary:

> NSA uses EO 12333 authority to collect foreign intelligence from communications systems around the world. Due to the fragility of these sources, providing any significant detail outside of classified channels is damaging to national security.

The only sensible observation we can make at this point is that a leaked document on the use of malware by the N.S.A. seems to suggest that the EO 12333 governs untargeted malware, and that USSID 18 only restricts it once the N.S.A. specifically knows who it is targeting [4]. This would be consistent with our earlier argument that the 1978 FISA definition of 'installing a device' (cf. 1801(f)(4)

[4] The N.S.A. statement seems to illuminate that all surveillance operations, even domestic ones, that do not fall with the 1978 FISA definition are regulated by EO12333. In this paper, we focus on advanced network surveillance operations conducted from abroad, but how to exactly draw the line between FISA and EO12333 applicability is an important subject for further research.

FISA) does not cover the advanced network manipulations we present in Section 3.2.

2.3.2 Weak Legal Protections for Americans under EO 12333.

Section 1.1 of EO 12333 provides that electronic surveillance should consider U.S. persons rights. The details are further specified in the underlying documentation, in particular in the heavily redacted USSID 18. In the Washington Post, a former N.S.A. chief analyst provided some background on the interplay between FISA and EO 12333 [17]:

> "Look, NSA has platoons of lawyers, and their entire job is figuring out how to stay within the law and maximize collection by exploiting every loophole," he said. "It's fair to say the rules are less restrictive under Executive Order 12333 than they are under FISA," the Foreign Intelligence Surveillance Act.

In spite of the redactions in USSID 18, we can make several new contributions to our collective understanding how legal protection for U.S. persons indeed are less restrictive under EO 12333.

Intentionally targeting U.S. persons. Essentially, section 2.4 of EO 12333 establishes that electronic surveillance operations *that fall under the EO 12333 regime* and *do not fall under the FISA regime* may intentionally target U.S. persons, as long as they meet the conditions summed up in s. 4.1. USSID 18. Before looking at those specific conditions, we mention that it is striking that a central passage of the opening paragraph of section 4.1. is redacted. It reads:

> 4.1. Communications which are known to be to, from or about a U.S. PERSON [*one entire line redacted*] not be intentionally intercepted, or selected through the use of A SELECTION TERM, except in the following instances:

Here we can only call attention to the redaction, but we have no possibility of knowing what it exactly states. (This would be one of many specific points that could be clarified through *e.g.,* political oversight or a FOIA request.) The more specific 'instances', where 'communications which are known to be to, from, or about U.S. persons' may be 'intentionally intercepted' are outlined in sections of 4.1.(a-d) USSID 18 and span 4 full pages.

Even with the many redactions, we can see that the restrictions provide less protection on critical points than the already permissive 'minimization procedures' under FISA. Often, instead of FISA Court approval, operations merely need Attorney-General or in some cases even only Director N.S.A. approval to be 'legal'. To name just one relevant example, s. 4.1.(c).(1) of USSID 18 holds that when U.S. persons (including U.S. corporations) *consent* to a surveillance operation, the approval of the Director of the N.S.A. may suffice to go ahead with a program. Indeed, May 2014 saw revelations on N.S.A.'s 'strategic partnerships' with several leading corporations in several routing sectors, which may point at

obtained 'consent'. Given aforementioned legal loopholes, it seems likely that the advanced network protocol manipulations we discuss in Section 3.2 fall within this category, as well as subsequently collected internet traffic. But several of the relevant criteria are, again, redacted (notably an entirely redacted s. 4.1.(b).(a). on 'international communications'), which prohibits us from establishing this with complete certainty [19]. Political pressure or perhaps FOIA request specifically targeted at transparency on s. 4.1. USSID 18 would be particularly useful to further analyze the scope of unilateral approval by the Director of the N.S.A. for broad surveillance programs.

Wide Exemptions to Process U.S. Person data Already Collected. Under USSID 18, further processing of foreign communications is unrestrained (cf. s. 5.3 USSID 18), while the exemptions to further process communication between U.S. persons intercepted during the collection of foreign communications are generous (cf. s. 5.4.(d) USSID 18): when communications are encrypted; when 'significant' for a 'foreign intelligence' purpose; when useful as evidence in criminal proceedings, or helpful to reveal communications security vulnerabilities (cf. section 5.4.d). In all these instances, the Director of the N.S.A. can determine to hold onto the communications between U.S. persons; under FISA, the Attorney-General must make such determinations.

Many More Classified Guidelines. Adding to the host of relevant but classified sentences, section 2 of USSID 18 references several classified legal documents that further govern specific intelligence activities, among them the DoD Directives mentioned at the start of this section and NSA/CSS Policy No. 1 to 23, titled "procedures governing NSA/CSS Activities that affect U.S. persons", as revised 29 May 2009. The public has no ability to analyse (the recent versions of) these mostly classified documents in detail. It is hard to tell whether such guidelines provide more legal loopholes. More generally, there are several differences in legal protection between FISA and EO 12333 that we haven't discussed here. In future work, we hope to further research these issues, in particular the DoD Directives, NSA/CSS policies and possible new disclosures.

2.3.3 Reforming the Third Legal Regime under EO 12333: Just One Branch of Government.

A more fundamental difference can be signalled at this point: over the next years, three branches of Government are involved with Patriot Act and FISA reform. In the sphere of EO 12333, this is not the case. Electronic surveillance regulated under EO 12333 is solely overseen by the Executive branch, regardless of its actual impact on the privacy of Americans. This simple observation has a long tradition in U.S. Constitutional law, that gives broad so-called Article II authorities to the U.S. President when it comes to national security. As we have seen, EO 12333 and its underlying guidelines have been adopted within the Executive Branch. Much of the lowered legal protection we have signaled demonstrates how oversight between branches of Government can be circumvented by conducting surveillance under EO 12333.

In response to media inquiries, Senator Dianne Feinstein (D-Cal), Chair of the U.S. Senate Intelligence Committee tasked to oversee U.S. intelligence agencies, provides some insight into what seems a complete lack of congressional oversight over EO 12333 operations [34]:

> "Twelve-triple-three [EO 12333] programs are under the executive branch entirely." Feinstein has also said the order has few, if any, privacy protections. "I don't think privacy protections are built into it," she said."

One very real outcome of this lack of oversight (or checks and balances *between* separate branches of Government) is the wide range of redactions still in place under the EO 12333 regime, which limits independent analysis. We mention, that by contrast, a central criterion under the European Convention of Human Rights ('ECHR') is that any government policy that impacts rights guaranteed under that Convention, including privacy, must be publicly available and contain specific safeguards against overreach. Under the ECHR, such transparency is seen to be a critical guarantee against overreach and abuse of power. In a range of cases, such as the 2008 Liberty v. UK case, surveillance programs have been ruled in violation of the fundamental right to privacy precisely because transparency and safeguards against abuse had been absent. [32].

To summarize, programs under EO 12333 may collect startling amounts of sensitive data on both foreigners and Americans. EO 12333 and USSID 18 presume communications are non-American, precisely because their operations are conducted abroad. Such operations are regulated by guidelines adopted almost entirely within the Executive branch, without any meaningful congressional of judiciary involvement. Generous exemptions exist that enable use of information 'incidentally' collected on U.S. persons, and critical details remain classified. Overcoming these concerns remains an issue that will be addressed entirely by the Executive branch. So far, it has not sufficiently been addressed at all, most probably because the lack of checks and balances between three branches of Government.

3 Loopholes that Exploit Network Protocols

We have just argued the that collection of US person's network traffic from abroad presents a major loophole that can be exploited to circumvent legal safeguards protecting Americans and oversight mechanisms in other branches of Government. Put differently, the current regulatory framework for network surveillance by intelligence agencies creates incentives for conducting surveillance on foreign soil, regardless of whether it actually affects American communications or not.

We now discuss how the technical details of Internet's core protocols can cause traffic sent by Americans to be routed abroad, where it can be collected under the most permissive third legal regime for network surveillance. We distinguish two settings: (1) situations where the vagaries of Internet protocols cause

Americans' traffic to *naturally* be routed abroad, and (2) situations where Internet protocols can be *deliberately manipulated* to cause Americans' traffic to be routed abroad.

3.1 Why US Traffic can Naturally be Routed Abroad.

The Internet was not designed around geopolitical borders; instead, its design reflects a focus on providing robust and reliable communications while, at the same time, minimizing cost. For this reason, network traffic between two endpoints located on US soil can sometimes be routed outside the US.

3.1.1 Interception in the Intradomain.

A network owned by a single organization (even an organization that is nominally "based" in the U.S. such as Yahoo! or Google) can be physically located in multiple jurisdictions. The revealed MUSCULAR/TURMOIL program illustrates how the N.S.A. exploited this by presuming authority under EO 12333 to acquire traffic between Google and Yahoo! servers located on foreign territory, collecting up to 180 million user records per month, regardless of nationality [17].[5] Yahoo! and Google replicate data across multiple servers that periodically send data to each other, likely for the purpose of backup and synchronization. These servers are located in geographically diverse locations, likely to prevent valuable data from being lost in case of failures or errors in one location. The MUSCULAR/TURMOIL program collects the traffic sent between these servers: while this traffic can traverse multiple jurisdictions, it remains with the logical boundaries of the internal networks of Yahoo! and Google. Thus, we already have one example where loopholes under the legal regime of EO 12333 were exploited in the *intradomain*, *i.e.,* within the logical boundaries of a network owned by a single organization.

3.1.2 Interception in the Interdomain.

Another possibility is the *interdomain* setting, where traffic traverses networks belonging to different organizations. Specifically, interdomain routing with BGP can naturally cause traffic originating in a U.S. network to be routed abroad, even when it is destined for a network that is located on U.S. soil.

BGP (*i.e.,* the *Border Gateway Protocol*) is the routing protocol that enables communication between networks owned by different organizations (*Autonomous Systems* or *ASes*, *e.g.,* Google's network, China Telecom's network, or Boston University's network). As shown in Fig. 1, ASes are interconnected, creating a

[5] The Washington Post has also revealed some aspects of a similar program called INCENSER (of which the technical details remain unknown), that apparently collected 14 *billion* user records in the same 30 day period. Together, these programs are grouped under a still-classified umbrella program under the name WINDSTOP. And according to journalists with access to the source material, many similar programs exist [16].

graph where nodes are ASes and edges are the links between them. ASes use BGP to learn paths through the AS-level graph; an AS discovers a path to a destination AS via BGP messages that it receives from each of its neighboring ASes. An AS then uses its local routing policies to choose a single most-preferred path to the destination AS from the set of paths it learned from its neighbors, and then forwards all traffic for the destination AS to the neighboring AS that announced the most-preferred path.

Importantly, the local policies used to determine route selection in BGP are typically agnostic to geopolitical considerations; path selection is often based on the price of forwarding traffic to the neighboring AS that announced the path, as well as on the number of ASes on the path announced by that neighbor. This means that it can sometimes be cheaper to forward traffic through a neighboring AS that is physically located in a different country, rather than one located in the same country; this situation is common, for example, in South America (where network paths between two South American endpoint ASes often cross undersea cables to Miami [24]) and Canada (where network paths between two Canadian endpoint ASes regularly traverse American ASes [10]). Ongoing work by one of the authors seeks to measure how often this occurs when both endpoints are located in the US.

3.2 How Deliberate Manipulations can Divert US traffic Abroad.

In addition to situations where Americans' traffic is naturally routed abroad, the Internet's core protocols – BGP and DNS – can be *deliberately manipulated* to force traffic originating and terminating in an American network to be routed abroad. As we discussed earlier, deliberately manipulating Internet protocols for subsequent data collection from abroad, even when the manipulation was performed from within the U.S., does not fall under the legal definition for 'electronic surveillance' in FISA; instead, these manipulations are regulated under the most permissive third legal regime for network surveillance, EO 12333 (and perhaps further specified in non-public guidelines).

3.2.1 Deliberate BGP Manipulations.

We know of numerous real-world incidents where manipulations of the BGP protocol have caused network traffic to take unusual paths, including situations where traffic from two American endpoint ASes was rerouted through ASes physically located abroad. While there is no evidence that these incidents were part of a surveillance operation, or even a clear understanding of why they occurred, it is instructive to consider them as examples of how an authority could circumvent the legal safeguards protecting U.S. persons by forcing their network traffic to be diverted abroad.

In 2013, Renesys observed a number of highly-targeted manipulations of BGP that caused traffic sent between two American endpoint ASes to be routed through Iceland [26]. One manipulation that occurred on June 31, 2013, is shown in Fig. 1. Traffic originating at an endpoint physically located in Denver and

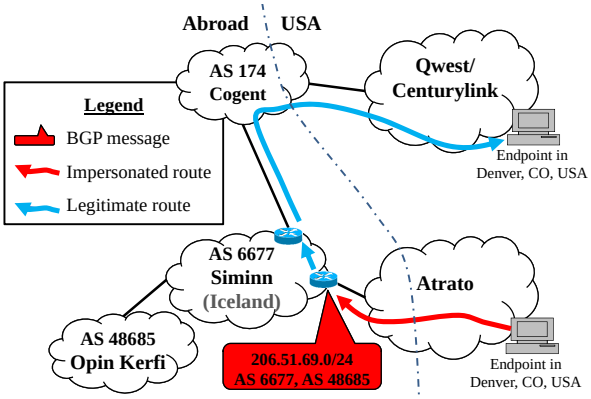

Fig. 1: On June 31, 2013, manipulator AS Siminn in Iceland used BGP to send an "impersonated route" for IP address block 206.51.69.0/24, allowing Siminn to intercept traffic sent between two endpoints in Denver, CO, USA. This incident was reported on by D. Madory at Renesys. [26].

logically located inside Atrato's AS, then travels to an Icelandic AS (Siminn) and then back to its destination, which is physically located in Denver and logically located in Qwest/Centurylink's AS. Renesys also observed nine other Icelandic ASes, as well as a few ASes based in Belarus, performing similar BGP manipulations.

Similar incidents have been known to occur periodically in the Internet [9]. In 2010, for example, a routing incident caused traffic sent between multiple American endpoint ASes to be diverted through China Telecom during a single 18-minute time period [11]. In 2008, a presentation at DEFCON [27] demonstrated how these manipulations could be performed in a covert manner that could be used to confound the network measurement mechanisms (*e.g.,* traceroute, BGP looking glasses) that researchers used to detect the 2010 and 2013 incidents mentioned above.

Target and Location of the BGP Manipulation. To understand how the legal framework applies to manipulations of the BGP protocol for the purpose of surveillance, we need to understand *who* is targeted, and *where* the manipulation is executed.

The incidents mentioned above are executed as follows. Per Fig. 1, the manipulating AS (*e.g.,* Icelandic AS Siminn) manages to divert traffic to itself by sending, to some carefully selected neighboring ASes, BGP messages that "impersonate" those sent by the legitimate destination AS (Qwest/Centurylink's AS). Because BGP lacks authentication mechanisms, these neighbors (Atrato's AS) accept the BGP message for the impersonated route,and select the impersonated route. They then forwards their traffic along the impersonated route to the manipulator's AS (Icelandic AS Siminn). The manipulator receives the traffic, and forwards it back to the legitimate destination AS (Qwest/Centurylink) via a le-

gitimate route. The manipulator AS therefore becomes a *man-in-the-middle* between targeted source AS (Atrato) and the destination AS (Qwest/Centurylink). While Fig. 1 shows traffic between two individual endpoints within Atrato and Qwest/Centurylink being intercepted by the BGP manipulation, typically *all* traffic originating inside Atrato and destined to the Qwest/Centurylink AS would be intercepted by the manipulator.

To further understand the targets of this manipulation, we consider what it means to send BGP messages that "impersonate" a legitimate destination AS. First, we provide more detail on BGP messages. A BGP message is used to advertise the path to a specific *IP address block* hosted by a particular destination AS.[6] Each AS in the Internet is allocated one of more IP address blocks, used to identify devices operated by that AS. Multiple devices can use a single IP address; thus, referring back to our legal analysis, a single IP address can be used by multiple devices or even 'persons'. A separate BGP message is used to advertise each IP address block allocated to a particular destination AS.

Thus, sending a BGP message that "impersonates" a legitimate destination AS means that the manipulator AS (Icelandic AS Siminn) sends a BGP message that claims a false route to the IP address block (206.51.69.0/24). As shown in Fig. 1, the manipulator AS (Siminn) falsely claims that the IP address block 206.51.69.0/24 is allocated to Siminn's own customer AS, the Icelandic Opin Kerfi AS 48685; in reality that IP address block is allocated to the legitimate destination AS (Qwest/Centurylink). Because BGP lacks mechanisms that can authenticate allocations of IP address blocks, the manipulator's neighbors will accept this impersonated route, and forward all traffic destined to the IP addresses in the disputed block to manipulator's AS (Siminn), instead of the legitimate destination (Qwest/Centurylink). This "impersonated" route will continue to propagate through the network, as the ASes that select the "impersonated" route pass it on to their own neighbors.

Thus, we can see that the "target" of this BGP manipulation is (1) all traffic sent by each source AS that selected the impersonated route (*e.g.,* all traffic from Atrato) that (2) is sent to IP addresses in the block that the manipulator falsely claims is allocated to him (*e.g.,* the 256 IP addresses contained in the block 206.51.69.0/24). That has important legal implications: the permissive legal regime under EO 12333 applies to such surveillance operations, as it does not necessarily 'intentionally' target a 'known, particular U.S. person'. It is also important to note that this BGP manipulation (which involves sending just a *single* "impersonated" BGP message from the Icelandic AS Siminn, shown in red in Figure 1) is executed entirely outside of the targeted endpoint ASes (Atrato

[6] An Internet Protocol (IP) address is a numerical address used to identify a particular device connected to the Internet; IP addresses are 32-bit numbers, divided into four 8-bit octets (written as *e.g.,* 206.51.69.201). An IP address block is a set of IP addresses that have a common n-bit prefix. For example, the set of IP addresses {206.51.69.0, 206.51.69.1, ..., 206.51.69.255 } has a common 24-bit prefix. We write this as address block 206.51.69.0/24, where the notation /24 ("slash twenty four") implies a common 24-bit prefix (here 206.51.69) for all addresses in the block.

and Qwest/Centurylink). Thus, this BGP manipulation *can be executed entirely abroad*.

3.2.2 Deliberate DNS Manipulations.

An alternate network protocol manipulation that can divert traffic to servers located abroad involves manipulating the DNS (*i.e., Domain Name System*). The DNS is a core Internet protocol that maps human-readable domain names (*e.g.,* `www.facebook.com`) to the IP addresses that identify the servers hosting the domain (*e.g.,* 69.63.176.13); applications that wish to communicate with the domain (`www.facebook.com`) first perform a *DNS lookup* to learn the IP address of the server that hosts the domain, and then direct their network traffic to that IP address. DNS lookups for end users and applications within a single AS are typically performed by a device called a *recursive resolver*, typically located within the AS; see Fig. 2. Recursive resolvers engage in the DNS protocol with devices located outside their AS, and return responses to DNS lookups to users and applications within their AS.

The DNS is well known to be vulnerable to manipulations the subvert the mapping from a domain name to IP address [8, 20, 22].[7] These manipulations, which have often been observed in the wild as mechanisms for performing network censorship [1, 35], can also be used to redirect network traffic through servers located abroad. Fig. 2 presents an example. Suppose that a manipulator wants network traffic destined to `www.facebook.com` from a given source AS (*e.g.,* Boston University) to be routed though a foreign server located abroad. Suppose the foreign server has IP address 6.6.6.6. The manipulator can execute a DNS manipulation that causes the recursive resolver in the source AS (Boston University) to map `www.facebook.com` to IP address 6.6.6.6. All network traffic for `www.facebook.com` from the source AS (Boston University) will then flow to the foreign server at IP address 6.6.6.6. Finally, the foreign server will silently forward the traffic it receives to the real facebook server at IP address 69.63.176.13. Thus, the foreign server becomes a man-in-the-middle for traffic sent between two US endpoints (Boston University and `www.facebook.com`).

Target and Location of the DNS Manipulation. This manipulation is more fined-grained than the BGP manipulation we discussed earlier: it targets all traffic sent to a particular domain (`www.facebook.com`) that is sent by all users and applications served by the targeted recursive resolver (*i.e.,* within a Boston University's AS). Again, the permissive legal regime under EO 12333 applies to such surveillance operations, as the traffic does not necessarily 'intentionally target a U.S. person'.

Moreover, like the BGP manipulations we described earlier, *these DNS manipulations can be conducted entirely abroad*; Hertzberg and Shulman [20] describe a technique that allows this manipulation to be executed by a device

[7] Indeed, these vulnerabilities have motivated the development of DNSSEC, a security-enhanced version of DNS; however, DNSSEC is far from being fully deployed, so these vulnerabilities remain exploitable today.

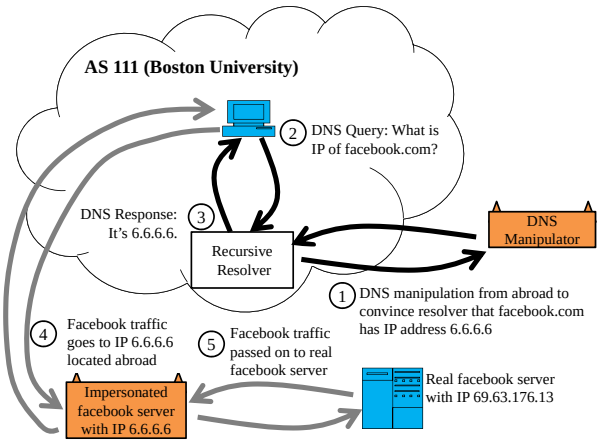

Fig. 2: Schematic showing how DNS manipulations can be used to direct traffic between two American endpoints (Boston University and facebook) to be routed abroad. The DNS manipulation technique labeled (1) is described in more detail in Figure 3.

located entirely outside the targeted source AS. For those interested in the details, we sketch out the technique below and in Fig. 3:

Figure 3: We show how the manipulator located abroad can subvert the DNS mapping for `www.facebook.com` at the target source AS 111 (Boston University). First, it is important to observe that recursive resolvers usually do not accept messages from senders outside their AS; however, mailservers do. (Mailservers are devices that provide email services for an AS. They therefore need to accept emails from outside the AS.) Thus, a manipulator located outside the target AS can use the mailserver to attack the recursive resolver. Specifically, the manipulator sends some carefully-crafted messages to a mailserver located inside the target AS. These messages act as a trigger for the mailserver to send DNS queries to the DNS resolver inside the AS; the DNS resolver accepts messages from the mailserver, because the mailserver is inside the AS. The recursive resolver then proceeds to resolve the mailserver's DNS queries. To do this, the recursive resolver sends DNS messages to other DNS servers outside the target AS. Finally, the manipulator responds to these DNS messages with carefully-crafted bogus DNS messages of its own; this allows the manipulator to subvert the recursive resolver's mapping from a domain name to an IP address. Observe that this manipulation just involves sending messages from outside the AS; no internal devices in the AS need to be compromised.

3.2.3 Other Manipulations.

The BGP and DNS manipulations we describe fall outside of the 'intentional acquisition' and the 'installation of a (..) device' subsection of the 'electronic

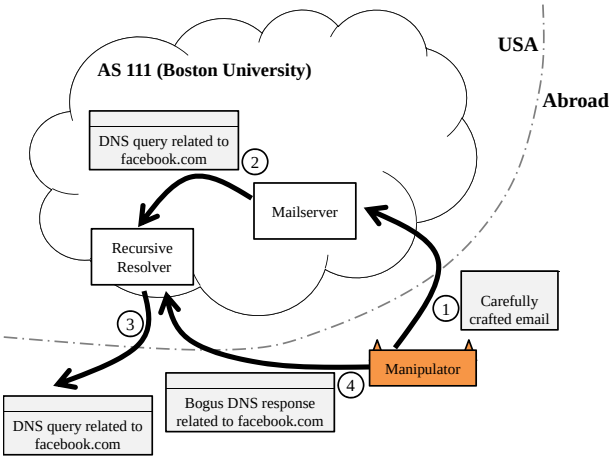

Fig. 3: Hertzberg and Shulman's [20] technique for subverting the DNS mapping for a particular domain (here, `www.facebook.com`) in a recursive resolver that serves a particular target AS (here, Boston University AS 111). The manipulator can be located entirely outside the target AS, and need only send DNS messages and emails. No devices within the target AS need to be compromised.

surveillance' definition under FISA. As noted earlier, such manipulation are therefore regulated by the permissive legal regime under EO 12333.[8] Even so, the regulations governing these network protocol manipulations under EO 12333 remain unclear, as many important legal documents remain classified or redacted.

For instance, FISA makes clear that protocol manipulations do not have to be executed entirely abroad to be regulated under EO 12333. To be completely confident that they can also be conducted on U.S. soil under EO 12333, one needs to have complete insight into USSID 18. On the face of it, EO 12333 and USSID do not define 'targeting' and FISA does not include manipulations within its scope.

While the BGP and DNS manipulations we described here can be executed entirely abroad, and regulated by EO 12333, there are whole other classes of manipulations that might be executed on U.S. soil. This class of manipulations includes any network exploit executed by an attacker that wishes to become a man-in-the-middle on a communication path. We will discuss more of these exploits in future work. Here we briefly mention a particularly interesting class of manipulation involving hacking into U.S. routers and switches and installing routes that deliberately cause traffic to be diverted abroad. Recent

[8] Note also that these FISA regulations were written in 1978, when the 'installation of a device' was perhaps necessary to divert traffic to a network location where it could be collected. Today, no installation of devices is necessary; instead, one can exploit vulnerabilities in already-present network devices (routers, web proxies, *etc.*) and network protocols (BGP, DNS, *etc.*) in order to alter the flow of network traffic.

revelations suggest that the N.S.A. does have the capability to take control of remote routers (*e.g.,* the HEADWATER, SCHOOLMONTANA, SIERRAMONTANA, and STUCCOMONTANA programs [21, 31]). We also know that the N.S.A. can *physically* tamper with U.S.-made routers [18]. Possibly also relevant is the N.S.A.'s SECONDDATE program, which the N.S.A. calls "an exploitation technique that takes advantage of web-based protocols and man-in-the-middle capabilities" [15]. At this point, however, we will not speculate as to whether or not the N.S.A.'s ability to subvert network protocols and routers is actually used to circumvent the statutory and constitutional protections provided to U.S. persons under the first two legal regimes we have described.

4 Summary and Future Work

In this paper, we highlighted a few of the major loopholes in the current legal regimes authorizing network surveillance by the U.S. intelligence community. International communications intercepted on U.S. soil are regulated by FISA and are subject to oversight by Congress and the judiciary. By contrast, surveillance on Americans from abroad under EO 12333 is by and large the sole domain of the Executive branch. Designing a surveillance operation to adhere to two main criteria — to not 'intentionally target a U.S. person' (like *e.g.,* bulk surveillance) and to be conducted abroad — allows the the operation to be regulated by the permissive legal regime under EO 12333, thus circumventing constitutional and statutory safeguards seeking to protect the privacy of Americans.

The legal loopholes we identified are exploitable, since the vagaries of Internet protocols can sometimes cause traffic sent between two US endpoints to be routed abroad. Even when this is not the case, core Internet protocols like BGP and DNS can be deliberately manipulated to ensure that traffic between US endpoints takes an unusual path through a device located abroad. If the two main criteria are met, these interdependent legal and technical loopholes enable largely-unrestrained surveillance on Americans communications. For instance, these techniques can be used to collect, in bulk, all communications sent from an autonomous system like Boston University to a given IP address block (with a BGP manipulation), or from an autonomous system to a particular domain like `www.facebook.com` (with a DNS manipulation).

In future work, we will consider additional technical loopholes, as well as legal and technical remedies that can address the difficulties highlighted by our analysis. We will discuss why technical solutions like encryption, DNSSEC, and the RPKI can help combat these risks, but still are no panacea. Even encrypted traffic, for example, exposes "metadata" (including who is communicating, the length of the communication, *etc.*); moreover, FISA and USSID 18 minimization procedures permit retention and analysis of encrypted communications even if two communicants are known to be U.S. persons. Meanwhile, the RPKI can limit the scope and impact of BGP manipulations, but does not completely eliminate them. Future work will also discuss possible solutions in the legal and policy space, including a more comprehensive analysis of the USA Freedom Act;

on the face of it, the proposed U.S.A. Freedom Act and 4th Amendment case-law concentrate on legal safeguards for U.S. persons, and offer little promise in closing the international surveillance loophole we have discussed here.

We reiterate that we do not intend to speculate on whether or not the intelligence community is exploiting the interdependent technical and legal loopholes that we have described. Instead, our aim is to broaden our understanding of the possibilities and deeper issues at hand. Indeed, our analysis has highlighted a central problem in law; namely, that law has an old-fashioned focus on physical materiality, in the sense that it matters much where surveillance is conducted. The networked communications environment challenges such conventional laws with a new technical reality, that does not respect the traditional geopolitical boundaries to which current constitutional and statutory protection are tailored.

Therefore, we emphasize that while the Patriot Act and FISA are overseen by all three branches of Government, EO 12333 remains solely under the executive branch because the U.S. Constitution grants it wide national security authorities to protect the nation against threats overseas. The implications for long term reform are real: even if the legislative of judiciary branches of Government address the loopholes in the Patriot Act and FISA, the U.S. Constitution emerges as a significant obstacle to the long term reform of EO 12333. We have argued that consolidation of the loopholes in EO 12333 within the Executive branch could leave Americans' Internet traffic as vulnerable to surveillance, and as unprotected by U.S. law, as the traffic of foreigners. Going forward, without a fundamental reconsideration of the lack of privacy and due process safeguards for non-U.S. persons, U.S. surveillance legislation leaves the door wide open for unrestrained surveillance on U.S. persons from abroad.

Acknowledgements.

We thank Ethan Heilman, Bruce Schneier, Haya Shulman and Marcy Wheeler for discussions and pointers that have greatly aided this work. Alexander Abdo, David Choffnes, Nico van Eijk, Ed Felten, Daniel K. Gillmore, Jennifer Rexford and the anonymous reviewers for HOTPETS'14 each provided insightful comments on drafts of this paper.

References

1. The open network initative. `http://opennet.net/`.
2. Exhibit a: Procedures used by the national security agency for targeting non-united states persons reasonably believed to be located outside the united states (...), as amended. `http://www.theguardian.com/world/interactive/2013/jun/20/exhibit-a-procedures-nsa-document`,, July 29 2009.
3. Exhibit b: Minimization procedures used by the national security agency in connection with acquisitions of foreign intelligence: Information pursuant to section 702 of the foreign intelligence surveillance act of 1978, as amended. `http://www.lawfareblog.com/2013/08/the-nsa-documents-part-vi-the-2011-minimization-procedures/`,, October 31 2011.

4. NSA-Dokumente: So knackt der Geheimdienst Internetkonten. *Speigel Online*, December 20, 2013. http://www.spiegel.de/fotostrecke/nsa-dokumente-so-knackt-der-geheimdienst-internetkonten-fotostrecke-105326-13.html.
5. N.S.A. memordandum: The national security agency: Missions, authorities, oversight and partnerships. http://www.lawfareblog.com/wp-content/uploads/2013/10/NSA-August-9-2013-Memorandum-on-Missions-Authorities-Oversight-and-Partnerships.pdf, 9 August 2013.
6. U.S. Supreme Court 568 U.S. *No. 111025*, February 26, 2013. http://www.supremecourt.gov/opinions/12pdf/11-1025_ihdj.pdf.
7. A. Abdo. The most important surveillance order we know almost nothing about. ACLU Blog, December 30 2013. https://www.aclu.org/blog/national-security/most-important-surveillance-order-we-know-almost-nothing-about.
8. S. M. Bellovin. Using the domain name system for system break-ins. In *Proceedings of 5th USENIX UNIX Security Symposium, USENIX Association, Berkeley, CA*, 1995.
9. K. Butler, T. Farley, P. McDaniel, and J. Rexford. A survey of BGP security issues and solutions. *Proceedings of the IEEE*, 2010.
10. A. Clement, N. Paterson, C. McCann, A. Gamba, J. Obar, and J. Stevenson. Ix maps. http://ixmaps.ca/.
11. J. Cowie. Rensys blog: China's 18-minute mystery. http://www.renesys.com/blog/2010/11/chinas-18-minute-mystery.shtml.
12. L. K. Donohue. Bulk metadata collection: Statutory and constitutional considerations. *Harvard Journal of Law and Public Policy, Forthcoming*, 2013.
13. L. K. Donohue. Section 702 and the collection of international telephone and internet content. *Available at SSRN*, 2014.
14. S. C. for International Security and Law. The national security agency at a crossroads, panel 5: The content collection controversy. https://strausscenter.org/strauss-news/nsa-conference-audio-and-video.html, April 8, 2014.
15. R. Gallagher and G. Greenwald. How the NSA Plans to Infect Millions of Computers with Malware. *The Intercept*, March 12, 2014. https://firstlook.org/theintercept/article/2014/03/12/nsa-plans-infect-millions-computers-malware/.
16. B. Gellman and M. DeLong. One month, hundreds of millions of records collected. *The Washington Post*, October 30 2013.
17. B. Gellman and A. Soltani. Nsa infiltrates links to yahoo, google data centers worldwide, snowden documents say. *Washington Post*, October 30 2013.
18. G. Greenwald. How the NSA tampers with US-made internet routers. *The Guardian*, May 12, 2014. http://www.theguardian.com/books/2014/may/12/glenn-greenwald-nsa-tampers-us-internet-routers-snowden?r.
19. G. Greenwald. new slide deck surrounding publication book 'No Place to Hide'. *glenngreenwald.net*, May 13, 2014. http://glenngreenwald.net/pdf/NoPlaceToHide-Documents-Compressed.pdf.
20. A. Herzberg and H. Shulman. Fragmentation considered poisonous, or: One-domain-to-rule-them-all.org. In *Communications and Network Security (CNS)*, pages 224–232. IEEE, 2013.
21. J. H. Jacob Appelbaum and C. Stcker. Shopping for spy gear: Catalog advertises NSA toolbox. *Der Speigel*, December 29, 2013. http://www.spiegel.de/international/world/

catalog-reveals-nsa-has-back-doors-for-numerous-devices-a-940994.html.
22. D. Kaminsky. Black ops 2008: Its the end of the cache as we know it. *Black Hat USA*, 2008.
23. R. Levinson-Waldman. What the government does with americans' data. Technical report, Brennan Center for Justice at New York University School of Law.
24. D. Madory. 'crecimiento' in latin america. Renesys blog, May 23 2013. http://www.renesys.com/2013/05/crecimiento-in-latin-america/.
25. A. Narayanan and V. Shmatikov. Myths and fallacies of personally identifiable information. *Communications of the ACM*, 53(6):24–26, 2010.
26. A. Peterson. Researchers say u.s. internet traffic was re-routed through belarus. that's a problem. *Washington Post*, November 20 2013.
27. A. Pilosov and T. Kapela. Stealing the internet. DEFCON, 2009.
28. J. Risen and E. Lichtblau. Bush lets u.s. spy on callers without courts. *New York Times*, December 16, 2005.
29. J. Sensenbrenner. The usa freedom act h.r. 3361/ s. 1599. http://sensenbrenner.house.gov/uploadedfiles/usafreedomact.pdf, October 29, 2013.
30. C. R. Service. Reauthorization of the fisa amendments act. *7-5700, R42725*, April 8, 2013.
31. D. Storm. 17 exploits the NSA uses to hack PCs, routers and servers for surveillance. *Computer World Blog*, January 3, 2014. http://blogs.computerworld.com/cybercrime-and-hacking/23347/17-exploits-nsa-uses-hack-pcs-routers-and-servers-surveillance.
32. J. Van Hoboken, A. Arnbak, and N. Van Eijk. Obscured by clouds, or how to address governmental access to cloud data from abroad. *PLSC 2013, Available at SSRN 2181534*, June, 2013.
33. J. Van Hoboken, A. Arnbak, and N. Van Eijk. Cloud computing in higher education and research institutions and the usa patriot act. *Available at SSRN 2181534*, November, 2012.
34. A. Watkins. Most of NSA's data collection authorized by order ronald reagan issued. McClatchyDC Blog, November 31 2013. http://www.mcclatchydc.com/2013/11/21/209167/most-of-nsas-data-collection-authorized.html#storylink=cpy.
35. J. Zittrain and B. Edelman. Internet filtering in china. *IEEE Internet Computing*, 7(2):70–77, 2003.