

Deployment of Source Address Validation by Network Operators: A Randomized Control Trial

Qasim Lone*, Alisa Frik†, Matthew Luckie‡, Maciej Korczyński§, Michel van Eeten*, Carlos Gañán*

*Delft University of Technology, The Netherlands

†ICSI, UC Berkeley, USA

‡University of Waikato, New Zealand

§Univ. Grenoble Alpes, CNRS, Grenoble INP, LIG, France

Abstract—IP spoofing, sending IP packets with a false source IP address, continues to be a primary attack vector for large-scale Denial of Service attacks. To combat spoofing, various interventions have been tried to increase the adoption of source address validation (SAV) among network operators. How can SAV deployment be increased? In this work, we conduct the first randomized control trial to measure the effectiveness of various notification mechanisms on SAV deployment.

We include new treatments using nudges and channels, previously untested in notification experiments. Our design reveals a painful reality that contrasts with earlier observational studies: none of the notification treatments significantly improved SAV deployment compared to the control group. We explore the reasons for these findings and report on a survey among operators to identify ways forward. A portion of the operators indicate that they do plan to deploy SAV and ask for better notification mechanisms, training, and support materials for SAV implementation.

I. INTRODUCTION

Distributed Denial of Service (DDoS) attacks remain a significant challenge for network operators. In a 2019 survey by RIPE NCC of more than 4,000 participants, operators identified DDoS as the most critical security problem [1]. Attacks keep increasing in size. In February 2020, Amazon web services received the largest DDoS attack observed to date, which peaked at approximately 2.3 Tbps and lasted three days [2]. IP spoofing—sending Internet Protocol (IP) packets with a false source IP address—continues to serve as a primary attack vector for large-scale DDoS attacks [3]. It is used in amplification attacks, where an attacker forges the victim’s IP address in requests sent to systems that act as amplifiers, such as DNS or Memcached servers. These systems reply with larger responses than the request sent by the attacker, thereby congesting victim’s network or server. IP spoofing is also used in SYN flood attacks, to obscure the origin of the attack traffic.

The scourge of IP spoofing has Internet Hall of Fame technologist Paul Vixie [4] to observe: ‘Nowhere in the basic architecture of the Internet is there a more hideous flaw than in the lack of enforcement of simple source-address validation (SAV) by most gateways.’ Over the last decade, a movement of sorts has emerged around a manifesto on routing security [5]. It aims to remediate this problem by encouraging network operators to adopt a best current practice referred to as BCP38 [6]. BCP38—also more generally referred to as SAV—defines a method for routers to validate the source

address of every outgoing packet. A router should drop packets if the source address is not valid for the attachment point. Around 25-32% of the Autonomous Systems (ASes) tested by volunteers of the Spoofer project are reported to have problematic or wholly lacking SAV adoption [7].

This brings us to our main question: How can more operators be moved to adopt SAV? Earlier work on other security issues found that operators do act on notifications that report vulnerabilities or abuse in their networks, albeit to varying degrees [8]–[11]. Specifically for SAV, researchers at the Spoofer project recently reported that notifying operators boosted remediation rates by about 50% [7]. Their findings were based only on observational data. The authors argued that “ideally” one would undertake A/B testing to more reliably measure the effect of various interventions on remediation.

In this paper, we present the first randomized control trial (RCT, also called ‘A/B test’) for measuring the impact of notifications sent to 2,320 network operators on SAV remediation rates. This population is much larger than in any prior study on SAV. It is possible because we use misconfigured open resolvers as vantage points [12], [13]—a different technique to observe the lack of SAV adoption than the volunteer-based Spoofer project [14]. We include a control group in the design, which no earlier study on SAV did and which yields a crucial insight that puts the earlier findings in a different light: the improvements that [7] observed might be incorrectly attributed to the interventions.

Our study is novel in other aspects as well. We contribute to the research on notification mechanisms by conducting the first test of social and reciprocity nudges in the message design. In terms of channels, we test private messages to operators versus notifying national CERTs versus using geographically-organized Network Operator Group (NOG) mailing lists. Sending notifications to a public forum (NOG) has not before been tested in an experiment. Finally, we partnered with NIC.br, a leading Brazilian CERT, to have them deliver the treatment first-hand. CERTs are a trusted partner in the operator community and a critical player in the security notification ecosystem, yet it has not been measured if their notifications have more impact than those of researchers or security companies. We complement our experiment by a survey among operators, to help us interpret the findings and identify ways forward. In short: we conduct the largest and

most rigorous study on improving SAV adoption to date, as well as advance the knowledge on notification mechanisms.

Our study reveals a painful and disappointing reality: there is no evidence of any remediation driven by any of the treatments, compared to the control group. This includes treatments that prior work has thought to be effective. Even for the notifications from the Brazilian CERT, the trusted entity, we found no effect compared to the control group. Importantly, we did observe some remediation across all groups, including the control group. It might explain why [7] did report an impact of their notifications. Since they had no control group, they could not see that the remediation they measured was not actually driven by the intervention. All in all, our findings are sobering but important, if we are to correct our understanding of these interventions and move forward on this critical issue. Our survey among operators helps us identify how. In sum, we make the following contributions:

- We present the first rigorous notification experiment with a control group that focused on network operators as the primary population to be incentivized to adopt more secure practices.
- We perform the first large-scale notification experiment to measure the impact of social and reciprocity nudges in the notification messages, the use of a public forum (NOG mailing lists), and a national CERT sending out the notifications. None of the treatments performed better than the control.
- We use a Cox mixed-effects model [15] to quantify the impact of network complexity factors and socio-technical country level effects on the deployment of SAV. Our results show that smaller networks with fewer edge routers are likely to implement SAV faster than larger networks.
- Our survey confirmed that notifying contacts registered in WHOIS does, in fact, mostly reach the operator staff responsible for implementing SAV. The reasons they give for not implementing it are a lack of time and technical expertise. About half of respondents do indicate that they plan to implement SAV in the future. To improve SAV adoption, the operators recommend better notification systems, training on SAV implementation and better supporting resources like software and technical documentation.

II. RELATED WORK

In this section we review the existing methods to infer the adoption of SAV among network operators, prior experiments that tested the effectiveness of security notifications, and literature on nudging.

A. Methods to Infer the Adoption of SAV

Previous work [7], [14], [16]–[23] have proposed methods to detect networks that do or do not implement the SAV standard. They differ with respect to the direction of filtering, whether they infer the presence or absence of SAV, and

whether the measurements can be performed remotely or from inside the network under test.

The Spoofer project [7], [14], [20], [21] develops and supports a client-server system based on volunteers that run the client software from inside their networks. The client periodically sends and receives packets with spoofed source IP addresses to test if the SAV is deployed for both inbound and outbound traffic.

Lone et al. [23] described a remote method that relies on traceroute loops. When a packet is sent to a destination network with a routable but unallocated IP address space, it is forwarded back to the provider router, thus resulting in a loop. Such a packet should be dropped by the provider router as the source IP does not belong to the customer network. The main limitation is that it relies on a routing misconfiguration, and therefore coverage of the method is relatively small.

Müller et al. [16] and Lichtblau et al. [17] passively analyzed inter-domain traffic at large inter-connection points (IXPs) to detect networks not deploying SAV. However, the proposed methods need to overcome several challenges to be effective, such as analyzing noisy BGP data sources, AS relationship inference, and require collaboration with IXPs.

To detect the lack of SAV for outbound traffic, we implement a different method that does not require volunteers for vantage points inside the tested network and that enabled us to include a larger sample of operators in our study than prior work.

B. Security Notification Experiments

There has been a rich stream of studies on the effectiveness of notifications to operators of networks, websites and DNS infrastructure. Cetin et al. [24] described how ISPs notified and quarantined customers who were running devices that were vulnerable to being abused in amplification DDoS attacks. They reported the quarantined users achieved very high remediation rates, around 87%, even though these devices did not pose a risk to the users themselves and the users could easily exit the quarantine.

In another study, Kühner et al. [22] sent notifications to the network operators about open resolvers, which provide amplification and redirection for DDoS attacks in their network. They were able to remediate 92% of the open NTP servers which supported `monlist`. They used various intermediaries, including national CERTs, Network Operation Centers (NOCs), and notifications using the open NTP project. They, however, did not compare the effectiveness of these channels.

Luckie et al. notified network operators who had not implemented SAV in their networks [7]. They initially contacted them directly using email addresses listed in the WHOIS. Subsequently, they sent monthly emails to NOGs identifying ASes in a given region with apparent gaps in SAV deployment. They observed around 48.2% of remediation took at least 1 month to deploy. Furthermore, they reported NOG was twice as effective as private notifications.

Our work comes closest to the study by Luckie et al. [7], since we also notify network operators who have not imple-

mented SAV in their networks. However, their analysis of the impact of their interventions was based on observational data, not a randomized control trial. The lack of randomization and a control group makes causal inferences about the impact of notification on SAV deployment less reliable. Moreover, we also tested the significance of interventions using nudges and the impact of sending notifications through national CERTs on remediations. Another difference is that our study is based on a much larger sample [12], [13]. Our technique to detect SAV via open resolvers has two advantages over Spoofer data [7]: we find 10 times more providers that are not compliant (Oct 2020–Feb 2021) and we are not dependent on volunteers, so we can reliably re-check the identified networks for remediation.

In summary, we present the first study using randomized control trial to measure effectiveness of notifications on SAV remediations. We also tested the impact of social and reciprocity nudges on compliance.

a) Notification Channels: Previous studies have utilized various channels for reaching out to the network operators: the “abuse email” listed in the WHOIS database [25], physical letters [8], [26], and manually collected email addresses, postal addresses, phone numbers, and social media contacts [26]. Other studies used the authorized intermediaries, such as national CERTs [11], [12], [27], or clearinghouses, to deliver the notifications.

Max et al. [8] more than doubled the remediation rates for non-GDPR compliant websites (from 33.9% to 76.3%) by sending using physical letters instead of emails. Despite the effectiveness, sending notifications via post costs time and money: Maass et al. [8] spent around 5,000 Euros on postage alone to notify 3,997 non GDPR compliant websites. On the other hand, sending email using WHOIS record also presents challenges. Previous studies have experienced a bounce rate of over 50% in some cases [28], [29]. In our paper we prioritize contacts from `peeringDB` over WHOIS, where available. We explain this further in the methodology section.

C. Behavioral Nudges

Behavioral science literature suggests that nudges and minor changes in the framing of a message may lead to a higher compliance with a recommendation and drive the behavior change [30]–[32]. For example, in the security domain, previous studies have found that nudges are effective in motivating users to choose stronger passwords [33], update software [34], and make better online privacy and security choices [35]. Some common nudges utilize social comparison, authority, and reciprocity mechanisms to influence behavior. Specifically, *social comparison* raises normative behavioral expectations by contrasting target individual’s behavior with the behavior of other people in their social group [36], [37]. Making a request on behalf of *authority* is another persuasion technique leading to higher level of compliance than requests made by someone without authoritarian power [38]–[40]. Finally, in social psychology, *reciprocity* indicates a social norm that encourages people to respond to a positive or kind action with another positive or kind action [41]–[44]. For example, in

the ‘repeated helping game’ participants were more likely to provide costly help to other participants if they had received such help from them in previous rounds [45], [46].

In our study, we leveraged social comparison, authority, and reciprocity mechanisms in attempt to improve the effectiveness of notifications and nudge network operators to deploy SAV.

III. METHODOLOGY

We first explain the forwarders-based method for identifying operators who did not implement SAV. We then describe the experimental treatments and random assignment method. Finally, we discuss the design of the post-RCT survey.

A. Vulnerability Discovery

To identify networks that do not implement BCP38, we leverage a technique that uses misbehaving forwarding open resolvers as vantage points. It was proposed by Mauch [47] and later implemented by Kühner et al. [22] and Lone et al. [13]. Figure 1 illustrates the idea of the method. A Scanner (controlled by us) with IP 192.0.2.32 sends a DNS query to a misbehaving DNS Forwarder (with IP 203.0.113.54) to resolve the randomly generated `random.example.com` subdomain (Figure 1a). When the Forwarder receives the DNS query, it does not rewrite the source IP with its IP before forwarding it to a Recursive Resolver (e.g., 8.8.8.8) located outside the network under test. If the network hosting the vantage point has not deployed SAV, the forwarded query will reach the Recursive Resolver (Figure 1a: 2nd packet). The recursive resolver will perform a query resolution and return the query response directly to the Scanner under our control. Another possibility is that when the Forwarder receives a DNS query, it correctly rewrites the source IP address with its IP address and then passes it to the Recursive Resolver (Figure 1b). However, the forwarder sends the response from the recursive resolver to our scanner without rewriting the source address (Figure 1b: 4th packet). If the network does not implement SAV at the network edge, it will arrive at our Scanner with a spoofed IP address belonging to the Recursive Resolver.

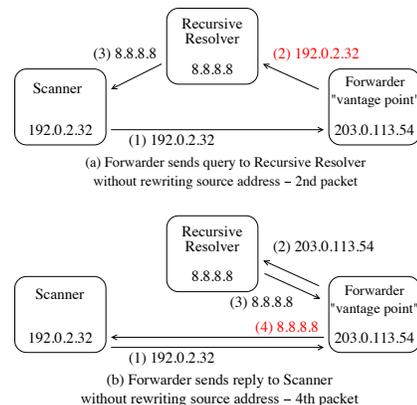


Fig. 1. Methodology to infer absence of SAV using forwarding resolvers.

We performed Internet-wide forwarders-based scans of IPv4 space weekly between September 2020 and February 2021 to identify misbehaving DNS resolvers in each routable network. We mapped their IP addresses to their ASNs and inferred 2,433 ASes operated by 2,320 providers in 118 countries had not at least partially deployed SAV for outbound spoofing. We also used the Maxmind GeoIP database [48] to map the IP address of misconfigured forwarders to their respective countries. Finally, we extracted contact addresses of the ASes using `peeringDB` [49] and `WHOIS` [25]. We also identified the relevant national CERT for each country using the APIs from FIRST [50] and SEI [51] and via manual search. The Spoofer project already sends notifications to NOG mailing lists. We utilized Spoofer’s NOG lists to map IP addresses in each country to the relevant NOG mailing list, if one was available.

The study population is network operators where we observed a lack of SAV with the technique explained above, which we operationalized as ASes with unique `WHOIS` contact email addresses. If two ASes had the same contact email address, we would assume they belong to the same operator and collate them. So to put it differently: the study population consists of 2320 unique `WHOIS` email addresses representing that number of operators.

Limitations of remediation tracking: Our data set that observes IP spoofing via misconfigured forwarders presents a few challenges to infer remediation. If a vantage point no longer shows up in our scan, this could mean the operator implemented SAV, but it could also mean the vantage point (temporarily) disappeared for other reasons. There could be DHCP churn [52], which means the forwarder’s IP address changed, though this will be to another address in the operator’s IP space. The user of the device could also switch it off. Or the operator could fix the misconfiguration, thereby making the device no longer send spoofed packets.

These factors mean that observations of spoofed traffic will appear and disappear also when there is no change in the adoption of SAV. If we have multiple vantage points for a network, then the impact of these measurement issues will be limited. Averaged across all weekly measurement cycles, we have more than one vantage point in 73% of all ASes. More importantly, the random assignment of our RCT design controls for this measurement problem. It will affect treatment groups and the control group more or less equally, meaning we can still reliably observe the impact of the treatments on remediation by looking at the difference among those groups.

To corroborate our findings on the presence or absence of SAV, we also included advice in the notification to run the Spoofer client, which can more directly observe SAV. However, only a small number of operators appeared to have done so (see Section IV-D). While Spoofer is more reliable, it requires volunteers to run the test and has lower coverage of networks than the open resolver-based method.

B. Experimental Design

To explore the effectiveness of notifications, we designed a large-scale randomized control trial (RCT) experiment. In an RCT, the subjects are randomly assigned to control and treatment groups. The effectiveness of the treatments are then assessed based on the comparison of the remediation rate in each treatment group with the control group. If the treatment is significantly different than the control group, researchers can confidently conclude that the intervention was successful.

We designed eight experimental treatments along two dimensions: delivery channels and message content. Figure 2 illustrates our experimental treatments, which we will now describe in more detail.

In every treatment group, using the communications channel associated with that treatment (see III-B1), we sent notifications about the discovered vulnerability and provided recommendations to deploy SAV, along with a link to the test that revealed the vulnerability and additional resources about remediation strategies. Beyond this baseline, in the nudging conditions, we added additional short nudging sentences (see III-B2). We also shortened the version of the baseline text for the NOG mailing list, to be consistent with the Spoofer notifications.

One of the requirements of randomized control trial experiments is to prevent contamination between the treatment and control groups. To fulfill this requirement, we built a public-facing website with private links for each operator with information only about their own network.

1) Notification Channel Treatments: We used three channels to deliver our notifications: (i) direct emails to the operators; (ii) emails to the national CERT, with the request to notify the non-compliant operators in their country, including Brazilian NIC; and (iii) emails to NOG mailing lists. In Brazil, we were fortunate to be able to partner with NIC.br, a trusted institution in a similar position as the national CERT. While NIC.br assured us to send the notifications to Brazilian operators assigned to the CERT treatment, we did not receive such assurance from CERTs in other countries. Therefore, NIC.br presents a special case within the CERT treatment group.

a) Direct Emails: The operators assigned to this treatment received the notification via a direct email. To find the contact addresses for ASes in our data set, we use the following process. We first check if there is a technical contact in either `peeringDB` [49] or `WHOIS` [25]. If both of them have an address and it is different, we prioritize the email address from `peeringDB`. We preferred `peeringDB` because it has been used in previous studies [53], [54] and they found the database up-to-date. If there are no technical addresses, we would use the listed abuse contact addresses, where we again prioritize the address from `peeringDB`. We preferred using the technical contact address, where possible, because we assumed that the odds would be higher to reach network engineering staff via that address rather than via the abuse address, which is managed by abuse handling departments.

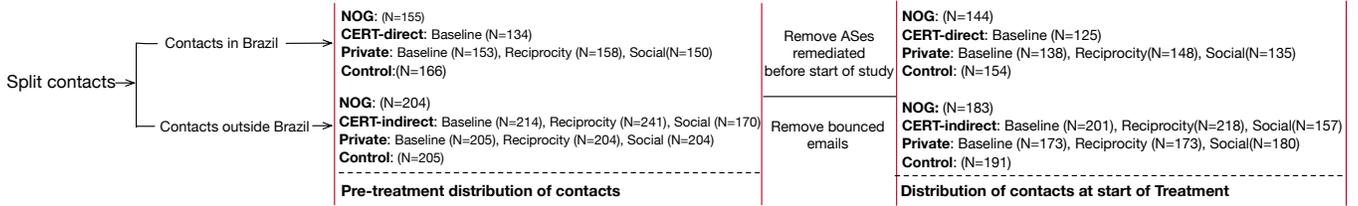


Fig. 2. Random assignment process and experimental treatments. The number of operators assigned to each treatment is included in parentheses

Implementing SAV requires reconfiguration of routers. This is better suited for the role of network administrators.

b) Notifications to CERTs: In the second treatment group, we sent the notifications to national CERTs and requested they forward the notifications to the operators. We asked CERTs to use the text of notification that we designed for the operators, to preserve the consistency of the notifications across groups (see Appendix A). Since this channel is indirect, it requires the cooperation of CERTs to forward our message to the relevant network operators. We have no way of ensuring that the messages were actually forwarded. This treatment leverages the CERT’s role and reputation (or *authority*, as discussed in Section II-C), so we can empirically measure whether they fulfill this role. We hypothesize that operators are more likely to take action if they receive a notification from CERT compared to an email from university researchers.

c) Notifications Directly from CERT: As we explained earlier, we partnered with NIC.br, a trusted CERT entity in the Brazilian operator community that routinely sends notifications about vulnerabilities to operators. This allowed us to set up a separate treatment where the CERT itself would issue the notifications. In contrast to the CERT treatment outside Brazil, the messages in the Brazilian CERT treatment would be directly sent by NIC.br, in Portuguese, and from their official email address. We hypothesized that the notifications are more likely to have impact if they come from an entity trusted by the network operator community. This allowed us to perform the first experimental test whether messages from CERTs, a critical player in the security ecosystem, have more impact than those of researchers. (An earlier study [27] also sent notifications to CERTs, but these were meant to be forwarded by the CERTs to the final recipients, the same as in our ‘notifications to CERT’ channel (b). The researchers could not ascertain if the CERTs actually forwarded the notifications.) To limit the effort required from NIC.br, we asked them to conduct only one treatment. This is consistent with how we approached all other CERTs: each received only a single treatment and a single message to forward to operators. Different CERTs were assigned to different treatments.

d) Notifications to NOGs: In the third group, we bundled our notification with the Spoofer notifications sent by the NOG lists. The Spoofer project measures the absence of SAV using a client-server application [7]. The project has been sending monthly emails since Dec 2018. The operators are used to these messages and already know that it is about missing SAV.

In terms of what operators are covered by either data set, the Spoofer data has minimal overlap with our open-resolver data. We discuss the comparison in more detail in the section IV-D.

The advantage of bundling the notifications and combining the measurements is that it saves network operators from receiving multiple emails about the same problem. Moreover, we hypothesize that publicly identifying the ASes on NOG mailing list would encourage them to deploy SAV more than when they receive this message through a private channel.

2) Nudging Treatments: In the CERT and private-email treatment groups, we differentiated our messages by incorporating specific nudges aiming at further motivating network operators to implement SAV. We created three conditions in each group: (i) the baseline message, which only contained the guidelines for the operators to understand the issue and how to fix it; (ii) the baseline message plus a social nudge; and (iii) the baseline message plus a reciprocity nudge. The full text of notifications is included in Appendix A.

In the social nudge condition, we urged the operators to deploy SAV and pointed out that most providers have already done so. To this purpose, we added following text to the content of the notification: “Note that 75% of network operators in the world already deploy BCP38 in their networks. Deploy BCP38 in your network to become one of them.”

In the reciprocity condition, we asked the providers to return the favor to operators who did implement SAV, thus reducing the attacks on everyone else, including the recipient. We added the following text to the baseline message: “Note that your network is receiving fewer DDoS attacks because other networks have deployed BCP38. Return the favor - deploy BCP38 in your network to make the Internet more secure.”

We chose encouraging (positive) framing of the nudges to the providers, rather than ‘naming and shaming’ (negative), because positive framing has been shown more effective in driving behavior change than negative framing [55]–[57].

3) Treatment Group Assignment: We use the data on the operators who lack SAV from October 2020 and randomly assign them—or more precisely, the unique WHOIS contact addresses for the ASes—to the experimental groups. We first separate the population in Brazilian and non-Brazilian operators (Figure 2). The special Brazilian CERT treatment meant we needed to randomly assign the Brazilian operators separately from the rest of the world. Then in both branches, we randomly assigned each operator to a treatment or control group. We had five treatment groups and one control group for the Brazilian sample, and seven treatment groups and one

control group outside of Brazil. The treatments are the same, except for the CERT group, which outside of Brazil includes two additional treatments for the social nudges. In total, we apply eight different treatments.

We had to modify the assignment process since CERT and NOG treatments operate at country-level: instead of assigning the operator contacts, we assign a country to a treatment group. The process becomes complicated since we want to have a balanced population across treatments, and the number of operators in each country is not the same. We designed our solution based on a best-effort algorithm to distribute contacts among different groups. We run the algorithm separately for contacts in Brazil and contacts in other countries. In each assignment, our random algorithm first validates if it can assign the contact to the treatment group. This is not always the case for the CERT and NOG treatments. For a few countries we have no contact point for a national CERT or for a NOG mailing list. If, for a specific operator, we have no CERT or NOG mailing list in our data set, the algorithm randomly selects another operator.

Under some conditions, the randomization could lead to unbalanced assignments. Stratification would then be used to ensure balanced treatment groups. However, methodological studies [58] have shown that in moderate and large samples, like ours, random assignment and stratification achieve similar variances. Furthermore, we checked various network and economic factors after the assignment to determine if the groups were in fact balanced. We statistically tested the group differences using ANOVA for: average AS size (i.e., number of IPv4 addresses calculated using longest matching prefixes in BGP announcement per AS), number of misconfigured forwarders, number of countries, number of stub ASes, membership of MANRS, Gross Domestic Product, and ICT Development Index assigned to each group. We found no statistical difference between the groups, which means they were similar for these variables.

4) Preventing Treatment Spillover: We designed the study to prevent contamination between the treatments. We built a website with an interface to the data on the non-compliant IP addresses and ASes. It also includes a detailed explanation of our methodology to infer the lack of SAV aided by dynamically generated diagrams containing misconfigured IP addresses and information on how to reproduce the result.

The website segments the information for different groups and recipients and does not contain any information for the control groups. We created separate sub-domains for CERT, NOG, and privately communicated treatments. We then generated unique URLs for each subject in the treatment. To prevent contamination within the private group, we sent individual links in our notification. The URLs only gave them access to the misconfigured IP addresses mapped to their ASes.

Similarly, we drafted a message for the CERT to forward to the ASes assigned to them. We instructed CERTs to append the AS number at the end of the URL to create a unique link for the operator they are contacting. Operators notified by CERT could potentially tinker with the URL to find information about

other operators assigned to the CERT group. However, they cannot find information about other treatment groups since a different sub-domain segregates them.

The notification to the NOG contains all the ASes and IP addresses assigned to the notified NOG. They cannot view operators assigned to other NOGs, since they are segregated via unique URLs. NOG treatment was likely to be seen by some operators in other treatments, but the NOG message had no information on operators in those other treatments. The website had no data on the control group.

C. Notification Procedure

We launched our first campaign on Oct 8, 2020 and sent notifications to 2,563 operators, and continued to conduct weekly scans to observe the remediation of IP spoofing. For operators that did not remediate, we sent a second message on Dec 8, 2020. We analyzed the remediation data until Feb 28, 2021. This meant that operators had about four months to implement SAV since our first notification.

Of all our emails, 102 (4%) bounced. In those cases, we retried with an alternate email address where possible, and reached additional 30 contacts. Eventually, we removed 72 contacts which we could not reach. Around 97% of our emails reached the recipients, which shows our approach to prioritize `peeringDB` and technical contacts gave improved reachability compared to previous studies, where in some cases the bounce rate was over 50% [28], [29]. In most cases, we got an automated reply that confirmed they had received the email and a ticket has been opened or someone would follow up. The German CERT copied us in cc in the forwarded notifications to the operators.

D. Post-Experiment Survey Design

To further understand the challenges in deploying SAV and contextualize the interpretations of our experimental findings, we designed a short survey aiming at collecting feedback from the operators. The survey has four main objectives. First, to understand the security challenges faced by network operators and what role SAV and DDoS play among them. Second, to understand if the notification has reached the correct contact person and preferable method for providers to receive similar notifications. Third, to understand the challenges in implementing SAV and whether the content of our notifications and referenced resources were sufficient for operators to deploy SAV in their network. Finally, we wanted to solicit suggestions on how to improve the notification process in general. Our survey was partially inspired by Lichtblau et al. [17], who in 2017 surveyed network operators about the impact of spoofing on their network, their filtering strategies, and challenges in the adoption of SAV.

In the survey, we asked participants about four main topics: 1) what security issues they believe their networks have, and how they discover them; 2) whether they have implemented or have planned to implement SAV and a subsequent question on their chosen methodology to deploy filtering from operators with SAV 3) who is responsible for implementing SAV in

their organization, and whether the issue was escalated to the responsible entity; 4) whether MANRS guidelines provide sufficient information on how to implement SAV, what other strategies can help achieve better compliance, and how would network operators prefer to be notified about IP spoofing issues. The full questionnaire is included in Appendix B. As compensation for their valuable time and comments, we offered all respondents a 50 Euro gift card through a raffle with a 1:15 chance of winning. Participants were offered an option to stay anonymous and let us donate the prize to a charity.

E. Ethics

We had a detailed discussion with the university’s IRB and received clearance to conduct the notification experiment and the survey. Our study followed all the active monitoring guidelines for ethical network measurement research [59], including creating a web page running at the IP address of the scanner, communicating with Internet response teams, and providing an opt-out option for operators.

We conducted our own scans since there is no existing public dataset that reveals non-compliance for SAV using our methodology. It is important to note that our scans are different from scans that aim to detect open resolvers, since we track responses that arrive from a different source IP address than the probed address. This means we cannot use existing data from open resolver scans conducted by Shadowserver and others. We randomly distributed our queries across the IPv4 address space, so the scanner does not consistently query the same AS before moving on to the next one. Furthermore, in line with the Menlo report [60], we considered that the marginal negative impacts of these measurements are outweighed by the beneficence of improved SAV adoption and reduced spoofed attack traffic. We only received two requests to opt-out and we immediately removed their IP ranges from the study.

Finally, we asked for consent from providers at the start of the survey and explained to them that we will anonymize their responses before publishing them. We offered compensation in the form of a lottery with gift cards. If they did not want to receive gift cards due to the nature of the job or for any other reason, we gave them an option to donate the amount to a charity and stay anonymous.

IV. RESULTS

In this section, we analyze the impact of our notifications on remediation rates across different treatment groups. We start by examining remediation at three different levels: organization, AS, and prefix level. Then, we compare the remediation rates between CERT and NOG treatments.

A. Organization-Level Remediation

We start the analysis at the organizational level. Organizations can operate multiple ASes and while the SAV compliance can differ per AS, the decision to implement SAV can be driven by organizational policies. Therefore, we bundled the ASes with the same contact email address together as they

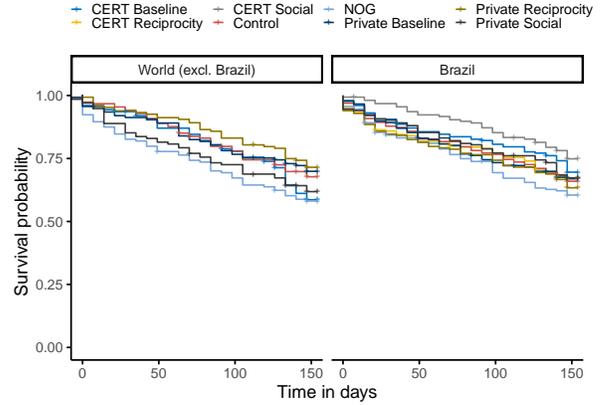


Fig. 3. Contact remediation survival plots for organizations in World excl. Brazil (left) and in Brazil (right).

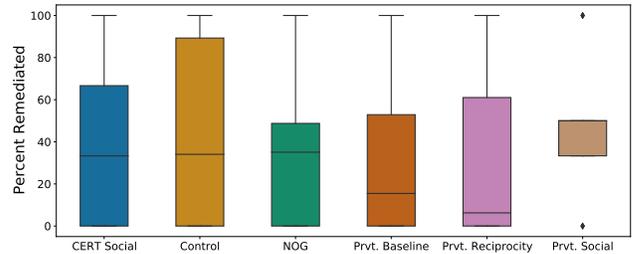


Fig. 4. Remediation per treatment group for countries that also received a notification in CERT Social group.

are most likely sibling ASes under the same administrative domain. Thus our unit of analysis is contact email addresses for the ASes. Our data set contains 200 (8.6%) contacts with more than one AS registered in WHOIS.

We only consider remediation as successful if all ASes under the contact email address do not appear in our scans after we have notified them. It is a high bar to pass since it might miss partial compliance, where providers might be remediating some ASes in their network or just a part of their AS.

To understand the differences across the groups, we compute the Kaplan-Meier survival curves per group as shown in Figure 3. On the y-axis, we have the probability of an organization deploying SAV t days after they received the notification (x-axis). This is estimated taking into account the number of organizations that had deployed SAV at time t divided by the total number of organizations that had not deployed SAV at time t . Overall, the survival curves show the same downwards trend for all the groups including the control. In Brazil, the NOG and Private Social groups do slightly better: they remediated 10% and 6% more than the control group, respectively. In the rest of the countries, networks in the NOG group remediated 5% more than in the control group.

To check whether these differences in remediation rates are statistically significant, we ran the log-rank test comparing the survival curves of the control group with the treatments. It

tests the null hypothesis $H_0 : S_1(t) = S_2(t)$ for all t where the two exposures have survival functions $S_1(t)$ and $S_2(t)$. We consider (≤ 0.05) as statistically significant. Confirming our initial visual observations, most of the groups did not have significantly different remediation rates. Only the result for the NOG group in Brazil is weakly statistically significant ($p = 0.049$). However, in light of how many treatments we tested, a 1 in 20 probability of this outcome being due to chance, is actually quite plausible. So we do not see this as enough evidence of an impact of that treatment group.

For all countries except Brazil, we also observed the CERT Social group remediated slightly slower ($p = 0.043$) than the control group. To understand why the CERT Social group remediated slower than the control group, we investigated the distribution of organizations at the start of our analysis in Figure 2. There are 34 (17.8%) fewer contacts in the CERT Social group than in the control. Hence, the baseline probability of remediation is also lower. Some network operators might have upgraded their routers or policies, which we count as baseline, or natural, remediation. In Figure 4, we compare remediation in the CERT Social group with other groups. We observe that remediation for contacts in the CERT Social group is similar to the control, NOG, and Private Social groups. Moreover, the average remediation in the CERT Social group is around 54%, while the average is only slightly higher for the rest of the countries (58%). In short, we can conclude that remediation in the CERT Social group is worse than in the control group mainly due to sampling differences.

B. Partial Remediation

An organization can choose to implement SAV for a few ASes but not for all the ASes they operate. Multiple ASes could also be managed by different teams, especially if these are located in different countries. Similarly, due to technical reasons like ASes not being stub or multihomed networks, operators might not be able to implement SAV in their entire network. To further investigate this, we analyzed partial remediation measured as the number of ASes and prefixes within an organization that implemented SAV within the study period.

AS-level remediation: Figure 5 shows the survival curves using ASes as unit of analysis. The results are almost identical to the organization-level results. The global remediation rates are not significantly different between the treatments and the control group. Only the ASes in NOG group in Brazil remediate significantly faster than the control group ($p = 0.05$).

Prefix-level remediation: Remediation can also occur at the prefix level, having both SAV compliant and non-compliant prefixes within the same AS. Figure 6 shows the survival curves of remediation using BGP prefixes as unit of analysis. Similar to both the organization-level and AS-level remediation, we observe no significant difference between the groups. Again, the only exception is the NOG group which remediated slightly faster than the rest of the groups.

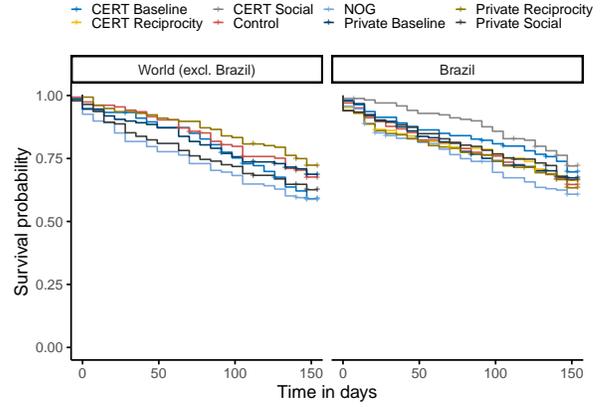


Fig. 5. AS remediation survival plots for ASes in the World excluding Brazil (left) and in Brazil (right).

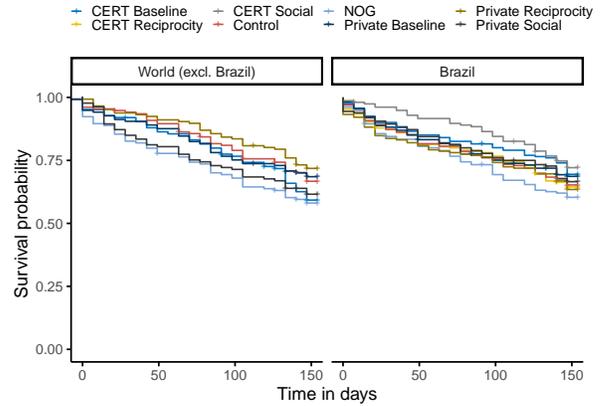


Fig. 6. Prefix remediation survival plots for World excluding Brazil (left) and in Brazil (right).

C. Main Experimental Effects

In this section, we analyze the differences in remediation rates across different experimental groups. We use relative risk ratio (RR) as a descriptive statistic to measure the probability of deploying SAV in one group compared to the probability of deploying SAV in the other group.

1) *Impact of the CERTs Groups:* We further compared the remediation across the CERT groups. Our motivation was to explore if there are significant differences between national CERTs. We calculate relative risk ratio between each pair of CERTs. In simple terms, this ratio produces a factor by which one CERT is different from the other in terms of remediation rate.

Figure 7 only displays the countries for which risk ratios—the differences in remediation—were significant. We determine the significance by looking at the confidence intervals (CI). If the CI includes the value 1, the RR is not statistically significant. If CI contain 1, it would mean that the relative remediations have no difference [61]. We interpret the figure row-wise for each national CERT. For instance, France had 4.2 times higher remediation rate than Argentina. In our sample,

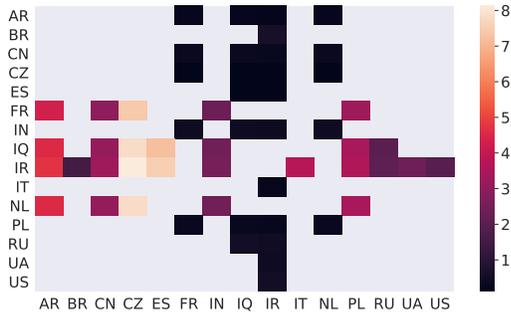


Fig. 7. Relative risk ratios among countries in the CERT group. Only the countries with significant risk ratios are displayed.



Fig. 8. Relative risk ratios among countries in the NOG group.

only networks in France, Iran, Iraq, and the Netherlands assigned to the CERT group were more likely to remediate than the other countries.

2) *Impact of the NOG Group:* We also calculated the relative risk ratios between the countries assigned to a NOG experimental group. Figure 8 only shows the countries that significantly differ in remediation. We used the confidence intervals to determine the significance as explained earlier. Germany, France and Lebanon NOG’s were more likely to remediate than other countries outside of Brazil in our sample. The RR for Brazilian NOG did not have any significant value, which in other words means that ASes in Brazilian NOG did not remediate more than other countries.

3) *Impact of Nudges on Remediation:* We explore the effectiveness of adding social and reciprocity nudges to the baseline text of notifications on remediation rates. We aggregate data for each of the nudging conditions (baseline, social, and reciprocity) from the different treatment groups and compare them against the control group. In Table I, we show the relative risk of remediation with reference to the control group. All of the nudges have a relative risk of around one compared to the control, which shows the nudges did not significantly impact remediation. In other words, operators that received the notification with a nudge were as likely to remediate as operators in the control group.

TABLE I
RISK RATIOS FOR NUDGING CONDITIONS COMPARED TO THE CONTROL GROUP

Group	Remediated	Exposed	RR	CI
Control	112	345	-	-
Baseline (no nudge)	206	637	0.99	0.82-1.2
Reciprocity nudge	175	539	1	0.82-1.22
Social nudge	150	472	0.97	0.80-1.19

D. Comparison with Spoofer

We requested operators to run the Spoofer tool [7] to validate if they have correctly deployed SAV. A total of 1,670 ASes submitted tests using the Spoofer tool in the study period (Oct 2019 - Feb 2021). While we cannot know if our request caused the operators to use Spoofer tool, the overlap between the ASes from the Spoofer tool and our methodology is around 12% (296 ASes). It signifies that our experiment did not get contaminated because of the Spoofer project. Note that the Spoofer project sends monthly notifications to NOG lists and often gets presented at conferences. MANRS also recommends using the Spoofer tool to test SAV deployment [62].

We also analyzed the remediations reported by the Spoofer tool [7]. In total across all Spoofer measurements, 98 ASes in Spoofer data implemented SAV in their network during our study period (Oct 2020 - Feb 2021). Of these, 22 ASes overlap with our measurements and 5 of them are in the control group. Since we did not send notifications to the control group, this clearly demonstrates that there is some natural remediation occurring. It is important to note that we sent notifications to 2,563 ASes which had not deployed SAV, while during the study period, the Spoofer dataset revealed only 248 ASes without SAV.

We can conclude from these results that there is limited evidence that operators acted upon our notifications. Moreover, positive remediation rates in the control group signals that factors other than our interventions influenced SAV as well.

V. FACTORS AFFECTING REMEDIATION RATES

Multiple factors could have affected remediation rates. Such factors could range from the size and complexity of the network, to the lack of budget and/or expertise. In this section we first identify potential factors that might have an impact on SAV implementation rates, and then quantify this impact through regression analysis.

In response to our notifications, three operators requested additional guidance or information. For instance, one operator claimed that his network was fully compliant. However, in further discussion, with the evidence from the measurements, he acknowledged that part of his network was recently upgraded and was not compliant. The operator subsequently implemented SAV in the network and did not reappear in our measurements. Other operators showed signs they lacked SAV knowledge. For example, two operators did not fully understand our measurement methodology and thought that we were notifying them about open-resolvers in their networks.

We responded with a detailed explanation of our methodology. We did not receive further responses.

This anecdotal evidence suggests that lack of information or knowledge could have influenced the operators’ decisions to not implement SAV in their networks. There could also be socio-technical reasons for non-compliance, such as operators in countries with low GDP based on Purchasing Power Parity (PPP), lower Internet penetration, and limited learning opportunities. To further understand the impact of these factors, we built a Cox proportional hazards model with mixed effects. At the multivariate level of analysis, we performed a two-level Cox proportional regression analysis to examine the effects of AS- and country-level characteristics on SAV implementation rate, and to determine the extent to which characteristics at the AS and country levels explain variations in SAV implementation rates. The multi-level Cox proportional hazards model allowed us to account for the hierarchical structure of the data. We hypothesize that ASes are nested within countries with different socio-economic characteristics. This suggests that ASes with similar characteristics can have different SAV implementation rates when operating in countries with different characteristics.

Using the multi-level Cox proportional hazards model, the probability of implementing SAV after receiving the notification was regarded as the hazard. We assessed the assumption of proportional hazard using visual inspections of graphs and statistical tests based on weighted Schoenfeld residuals. Two-sided p-values (≤ 0.05) indicated statistical significance. As explanatory variables we used socio-technical factors and a set of factors derived from the operators’ email responses, including the following:

CERT: boolean variable. `True` if the notification was sent to the national CERT, `False` otherwise.

NOG: boolean variable. `True` if the notification was sent to the NOG, `False` otherwise.

Private: boolean variable. `True` if the notification was sent to the technical contact email address of the AS, `False` otherwise.

AS size: numerical variable. We estimated the size of an AS by counting the number of advertised IPv4 addresses. We calculated the size using BGP data from Routeviews project [63]. We used weekly data for Oct 2020 and calculated the average IP space advertised by the ASes in our data set.

ISP: boolean variable. `True` if the AS belonged to an Internet Service Provider, `False` otherwise. To check whether an AS is used by an ISP we leveraged Telegeography: the GlobalComms database [64]. The database contains a highly reliable overview of the main broadband ISPs in each country, drawn from annual reports and market filings. The database contains details of major ISPs in 84 countries.

Edge Rtr: numerical variable. This variable is calculated by counting the number of edge routers of an AS. We used CAIDA’s Internet Topology Data Kit (ITDK) for March 2021 [65] to count the number of border routers per AS. The ITDK consists of routers and links observed in traceroute data collected from multiple vantage points, alias resolution

to identify which IP addresses belong to the same router [66], and a mapping from router to AS heuristically inferred using bdrmapIT [67]. We counted the number of border routers for ASes in our dataset connected to other ASes.

Stub: boolean variable. `True` if the AS is stub, `False` otherwise. We used Caida’s AS relationship data [68] to determine if the ASes in our data set are stub or not.

IDI: numerical variable. This variable represents the ICT Development Index (IDI) which is provided by ITU (United Nations International Telecommunication Union) and represents ICT development per country [69]. It assigns values from 1 to 10 to each country, with a higher value representing a higher level of development based on various ICT indicators.

In Table II, we present the results from the Cox model. The parameter estimates reported in the *est* column are log-hazard ratios. Their exponentiation produces hazard ratios. P-values indicate the statistical significance of each factor.

TABLE II
COX MIXED-EFFECTS MODEL WITH RANDOM EFFECTS FOR COUNTRIES.

Parameter	Est	Std.err	P-value	CI
<i>Fixed effects</i>				
CERT	-0.06	0.12	0.60	[-0.29; 0.16]
NOG	0.23	0.13	0.07	[-0.02; 0.48]
Private	-0.02	0.11	0.85	[-0.23; 0.19]
ASsize(ln)	-0.06	0.03	0.02	[-0.11; -0.01]
ISP	0.12	0.17	0.48	[-0.21; 0.44]
Edge Rtr(ln)	-0.05	0.02	0.00	[-0.08; -0.01]
Stubs	0.33	0.10	0.00	[0.13; 0.54]
IDI	-0.05	0.03	0.15	[-0.11; 0.02]
<i>Random effects</i>				
Group	Variable	Std Dev	Variance	
Countries	Intercept	0.217	0.04	

The notification channels did not impact significantly the implementation of SAV. Interestingly, only the NOG group has a positive coefficient which indicates that ASes that received a notification via this channel have higher probability of remediating than those in the control group. In particular, the hazard ratio for the NOG group is $\exp(0.23) = 1.25$. Therefore, notifying operators via NOG increases the probability of remediation by 1.25 times compared to ASes that received no notification.

Regarding the impact of AS size on SAV deployment, the argument can be made on both sides. For instance, bigger networks are more likely to have more resources to implement SAV. On the other hand, smaller networks are likely to have less complex networks and hence require relatively simpler configurations to implement SAV. In our results, we observe that smaller ASes were more likely to implement SAV in their networks. In particular, a 10% increase in the size of an AS, holding all other variables constant, was associated with a 5.82% decrease in the probability of SAV deployment.

The number of edge routers also decreases the probability of remediation. Network operators use multiple links to load-balance the traffic and avoid a single point of failure. To remediate, operators have to implement filtering policies near all edge routers. We found that networks with fewer edge routers were more likely to remediate after being notified. In

particular, a 10% increase in the number of edge routers in an AS, holding all other variables constant, was associated with a 4.87% decrease in the probability of SAV deployment.

There could be technical reasons preventing network operators from implementing SAV in their network. One factor could be having a non-stub or a transit AS. A customer of non-stub AS might not announce all routes to a provider because the AS is a customer of other providers as well. Hence, it is not technically feasible for provider ASes to apply strict filtering policies on their network [6]. We find that stub networks have 1.4 times higher remediation rate than the control group (holding all other variables constant). The country-level effect, an estimated intercept (excess risk) for each country, has a standard deviation of 0.21. This means that countries that are 1 standard deviation or more above the mean SAV remediation rate will have 1.24 times faster remediation rate than the norm, a modestly small country-level effect.

The other factors we considered did not significantly impact the remediation. One could hypothesize that ISPs would be more likely to implement SAV in their network since most end users are behind their networks and can be abused for an attack. While the hazard ratio sign indicates such relationship, we did not find statistically significant difference in remediation rate for networks that are ISPs compared to the control group. Finally, socio-economic factors defined by the ICT Development Index (IDI) did not influence the remediation, suggesting that the economic situation of a country has no impact on the remediation hazard.

In summary, we can conclude that network complexity plays an important role in remediation, i.e., the networks that are smaller in size and have fewer edge routers are likely to remediate faster. Similarly, stub networks are more likely to implement SAV faster in their network compared to non-stub.

VI. SURVEY RESULTS

To gain additional insights and feedback from the participants, we sent out the survey one month after our final notification. We sent a reminder to participate in the survey to non-responders after waiting for a month. We received responses from 32 network operators (less than 2%). While sample size does not allow us to make statistical comparisons between treatment groups, we believe that survey responses provide useful clarifications for interpreting our results.

a) Vulnerability Awareness: Ninety percent of survey respondents knew they had not deployed SAV, either because of the Spoofer tool test (30%), notifications from security researchers (20%), from NOGs (20%), from CERTs(10%), or based on their prior knowledge (10%). The remaining 10% were not sure if their networks deployed SAV.

b) SAV Implementation: Although 90% of respondents were aware that their network lacked SAV, more than half (52.7%) of the respondents reported that they have no filtering in place. Another 17% reported only partial implementation on some segments of the networks. Only 26% have implemented SAV throughout their network, and 4% were not sure.

More than half of respondents (53%) filtered out packets with a source IP address within private address space (RFC1918), so that only packets with a source address from routable IP space leave their network. It is important to note that filtering RFC1918 is simple as it has static address space and the filtering mechanisms are widely available. Lichtblau et al. [17] reported 70% of participants in their survey filtered RFC1918 addresses.

Moreover, 30% of respondents deployed SAV on routers that were customer-facing, 11% on their stub ASes, and 6% deployed SAV towards peering/IXP interfaces as well. In other words, they have deployed SAV in user space and those IPs cannot be abused to send spoofed traffic.

When we asked participants if they planned to deploy SAV in the future, we received mixed responses. Around 42% said that they were planning to deploy SAV, 33% had no plan, and 25% were not sure. One provider also sent us an email in response to our notification, saying that he acknowledges the issue and will get back to it after implementing another security practice (RPKI) in his network. Given that non-compliance is not an active “battleground,” it is likely that some providers assign SAV deployment to a lower priority compared to other network issues, but they might return to it later. However, we still think that 4 months we gave to the participants provided sufficient time to plan and remediate the issue, yet, we did not observe a significant impact on the outcome.

c) Notification Targets: It is possible that despite awareness, the respondents did not implement SAV, simply because they are not responsible for it. We wanted to confirm whether we reached the operator staff responsible for implementing SAV. There could be multiple reasons for not reaching the operator staff responsible for implementing SAV. For instance, 83% of the contacts we notified only had the address of the abuse mailbox. The abuse team is generally responsible for threats like spam, malware, and phishing campaigns from or towards the network. In cases where operators are not responsible, they may have another team performing network configurations.

However, a large majority (67%) of respondents said that they were responsible for implementing SAV. Only 13% said that they were not responsible, and 20% did not know what SAV is. Subsequently, respondents that believed they were not responsible for SAV said they did not escalate the issue to the responsible contact.

d) Reasons for Non-Compliance: We also asked operators why they had not implemented SAV in their networks. 30% of the respondents lacked the technical knowledge on how to perform filtering, and 30% lacked time to implement SAV. Another 18% were concerned that implementation may cause downtime or other performance issues. 12% mentioned technical reasons (multi-homed network, non-stub network) for not implementing SAV. Finally, 6% of the respondents thought that SAV is ineffective in addressing the attacks that use spoofed source addresses.

We can conclude from the survey results that the main reasons for non-compliance are driven by misaligned incentives and lack of knowledge, which are relatively easy to improve, compared to the concerns related to downtime, performance, or technical limitations.

e) Respondents' Suggestions for Improvements: In the final section of the survey, we asked participants for suggestions about possible improvements in the notification process. We sent MANRS guidelines [70] as part of our notification. About 73% of the respondents said that MANRS had sufficient information explaining how to implement SAV. However, 23% were not sure, and 4% said that MANRS does not provide sufficient details. They explained that the guide currently provides configurations only for CISCO and Juniper routers, and needs to cover configurations for other brands of routers as well. For example, one of our respondents said they used a Mikrotik router, which is not covered in MANRS.

One respondent suggested to create a dedicated channel for SAV notifications, where operators can also discuss technical difficulties in implementing SAV. 64% of the respondents requested more community-driven seminars that discuss SAV implementation. Finally, 36% of respondents suggested that routers should provide user friendly configurations to implement SAV.

While the sample size of our survey does not allow us to extensively generalize the results, it still provides valuable insights. We provide recommendations for improving the notification process and policies for SAV compliance in section VII.

VII. DISCUSSION & CONCLUSIONS

In this section, we interpret our results, discuss issues that might have played a role in low remediation, and present future avenues for improving both notifications and SAV adoption.

A. Treatment Effects

Except for the Brazilian NOG group, there are no significant differences when comparing remediation between the treatments and the control group. There can be multiple reasons why the Brazilian NOG group had higher remediation rates than the control group. First, operators that have subscribed to a NOG show their willingness to understand and discuss network challenges. Second, it creates peer pressure because the names of ASes are publicly available, while they can ignore the private communication. Finally, operators might trust the NOG channel, since the communication was part of the already known Spoofer project [7].

B. Remediation in the Control Group

We also observed remediation in the control group, where we did not send any notifications. There could be several reasons for that. First, some network operators might have upgraded their routers or policies, which we count as a natural remediation.

Second, some operators might have read articles or attended conference talks or seminars about current routing issues,

which could have urged the operators to adopt SAV. For instance, in the RIPE meeting in Oct 2020, with more than 1200 participants, MANRS presented their initiatives about routing security, including available resources to deploy SAV [71]. SAV is also discussed in various network operator conferences and channels, which might have further encouraged the adoption [72]–[74].

Finally, the MANRS program, which encourages members to be SAV compliant, has been very active in the recent years. They provide resources in the form of documentation, tutorials, and seminars to help network operators deploy best security practices. They reported that their members doubled in 2020, reaching 588 by the end of December 2020 [75].

While there can be many factors driving natural remediation, they affect all treatment and control groups equally. So we can still have confidence in our conclusions about the null effect of the treatments. This is the essence of the random assignment process: it neutralizes the impact of confounding factors.

C. Comparison with Previous Studies

Even though previous studies showed some success with large-scale notifications, our results show little to no impact. We attribute these to the following factors.

Complexity: Complexity can play a vital role in the success of notification studies. SAV requires significant time and expertise and can cause downtime if not correctly implemented. Previous studies ([27], [76], [77]) notified hosting providers and users about compromised websites which usually requires fixing the access privileges or removing malicious files. Similarly, other experiments [78], [79] notified web admins about misconfigurations or best practices for their domains. To properly configure their web server, the domain owners usually have to follow a set of simple steps in the notification. In comparison, SAV requires a thorough understanding of the network. The configurations and types of routers make it difficult to provide a similar guide. Finally, the recipients of the notification might need to escalate the issue to senior network operators since it requires downtime, and misconfiguration can cause major disruptions.

Target Audience: Multiple studies notified network operators about routing and security issues [7], [22]. However, none of these had a control group, which is required to reliably assess the effectiveness of remediation. Our study is the first one that focuses on network operators and performs a randomized control trial. Previous studies using RCTs either sent notifications to the domain owners [27], [76], [77] or to the network operator about compromised user devices [24]. In those cases, the operators are only asked to forward the message. They do not incur the main cost, as they rely on their users to remediate the problem.

Liability and incentives: The incentives of treatment subjects in our experiment is different from most operators of vulnerable or compromised resources. The benefits of implementing SAV flow to the rest of the Internet, not the operators themselves. The network implementing SAV is still vulnerable

to DDoS attacks from other networks. In terms of liability, a prior study had found higher remediation rates because of legal consequences [8]. However, there is no liability on operators prevent spoofed traffic from leaving their network.

Language of Notification: We sent out our treatments in English, except for those administered by the Brazilian CERT, which were in Portuguese. Notifications in network operators' native language could have improved the effectiveness of interventions. However, our study found no impact of the language difference. This is consistent with earlier work where more languages were included in a notification experiment, which also found no impact on remediation [11].

Awareness of Vulnerability: There has been a significant effort by the security community to deploy SAV over the last several years [7], [62], [80]. It is possible that some network operators already know through notifications from the Spoofer project that their network is non-compliant and have either ignored prior notifications or cannot deploy SAV due to technical limitations. That said, it is important to note that our dataset is very different from that used in the Spoofer-based campaigns, the main notification effort in this area. This dataset has not been used in previous notification campaigns.

D. Reasons for Non-Remediation

Our survey results found that 57% of respondents did not follow the recommendation to implement SAV, even though they confirmed we reached the right recipient in most cases. It contradicts previous work [17], where only 24% of the operators mentioned that they did not implement SAV in their networks. One possible explanation is that Lichtblau et al. [17] contacted only NOG members. The operators who have subscribed to the list are likely more aware of security challenges and willing to adopt best practices.

Our survey results revealed several reasons for non-compliance. Perhaps surprisingly, awareness about IP spoofing and the absence of responsibility for router configurations are not the prominent reasons. The majority of our survey respondents said that they were aware of the issue and were responsible for its remediation. Yet, many participants acknowledged that they were not familiar with how to perform filtering. Thus, as we discuss in Section VII-E, educating network operators about security vulnerabilities and remedies are important to improve compliance.

A large proportion of participants also mentioned that they lack time for implementing SAV, or that it is not their top priority. Finally, some respondents acknowledged concerns about performance issues or technical limitations deferring them from implementing SAV in their networks. While understanding relative impact of those reasons on remediation requires future work, our research and previous studies [1], [17] conclude that there is a need for community-driven efforts in aligning operators' incentives and providing better resources for addressing technical challenges with SAV implementation. We further discuss the recommendations for improving SAV adoption in the next section.

E. Moving Forward: Recommendations

Although notifications did not dramatically increase SAV adoption, we propose a number of steps that can help improve the adoption of routing and security vulnerability remediation.

Improving Notification Channels: Our survey response indicate that most of our notifications reached the recipients. However, to make sure they reach the team responsible for security and routing, we propose that providers should be encouraged by RIRs to fill in and keep up-to-date the technical team's contact details, in addition to abuse-email contacts.

Improving Resources: MANRS provides guidelines to network providers that describe how to implement SAV in their network, in English. To increase SAV adoption, it should be available in other languages, and it should cover other popular brands of routers in addition to CISCO and Juniper.

Improving Incentives: The main issue with routing security is that the remediation entails financial costs and requires human resources, while benefits would be mostly absorbed by the rest of the Internet. To align the incentives, the Internet community can play its part. Most of the providers with stub networks get connectivity through upstream providers. They hold a unique vantage point where they can detect if the incoming packets have a spoofed source [7], [17], [23]. If they exercise their position of power and peer with compliant networks, the overall compliance could increase significantly. There are examples where network providers leveraged their power to achieve compliance. For instance, a provider dropped invalid prefixes from its customer ASes [81]. The owners of the prefixes took corrective action and updated their Route Origin Authorizations (ROA) to fix the issue. Similarly, after observing a consistent BGP hijack from Bitcanal, Hurricane Electric and Portugal's IP Telecom were able to cut them off from the Internet [82]. Thus, the network community needs to take corrective actions. This could be supported by legislation that makes the providers liable for network attacks. Interestingly, two countries—Albania and the Philippines—consider avoiding correcting security flaws as administrative and criminal offenses [83]. Both inside and outside the network community, actions are possible to improve the incentives for SAV adoption.

ACKNOWLEDGMENTS

We thank NIC.br for assistance with this research, especially Gilberto Zorello, who generously translated notifications into Portuguese, sent them to the operators, and handled follow-up questions. The Spoofer project is the result of funding provided by U.S. DHS S&T Contract 140D7018C0010 and NSF OAC-1724853. This work was also partially supported by a grant from the Center for Long-Term Cybersecurity (CLTC) at U.C. Berkeley, by National Science Foundation grants CNS-1514211 and CNS-1528070, by the National Security Agency's Science of Security program. The published material represents the position of the author(s) and not necessarily that of funding agencies.

REFERENCES

- [1] RIPE NCC, “Survey Results,” 2019, <https://ripe79.ripe.net/presentations/89-RIPE-NCC-Survey-2019-Report-Presentation.pdf>.
- [2] “Amazon ‘thwarts largest ever DDoS cyber-attack’,” 2020, <https://www.bbc.com/news/technology-53093611>.
- [3] “NET SCOUT THREAT INTELLIGENCE REPORT,” 2020, https://www.netscout.com/sites/default/files/2020-02/SECR_001_EN-2001_Web.pdf.
- [4] P. Vixie, “Rate-limiting state,” *Communications of the ACM*, vol. 57, no. 4, pp. 40–43, 2014.
- [5] “MANRS,” 2020, <https://www.manrs.org/isps/participants/>.
- [6] F. Baker and P. Savola, “Rfc3704: Ingress filtering for multihomed networks,” 2004.
- [7] M. Luckie, R. Beverly, R. Koga, K. Keys, J. A. Kroll, and k. claffy, “Network hygiene, incentives, and regulation: Deployment of source address validation in the internet,” in *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, 2019, pp. 465–480.
- [8] M. Maass, A. Stöver, H. Pridöhl, S. Bretthauer, D. Herrmann, M. Hollück, and I. Spiecker, “Effective notification campaigns on the web: A matter of trust, framing, and support,” in *30th USENIX Security Symposium, USENIX Security 21*, 2021.
- [9] M. Carvalho, J. DeMott, R. Ford, and D. A. Wheeler, “Heartbleed 101,” *IEEE Security & Privacy*, vol. 12, no. 4, pp. 63–67, 2014.
- [10] Z. Durumeric, F. Li, J. Kasten, J. Amann, J. Beekman, M. Payer, N. Weaver, D. Adrian, V. Paxson, M. Bailey *et al.*, “The matter of heartbleed,” in *Proceedings of the 2014 conference on internet measurement conference*, 2014, pp. 475–488.
- [11] F. Li, Z. Durumeric, J. Czyz, M. Karami, M. Bailey, D. McCoy, S. Savage, and V. Paxson, “You’ve got vulnerability: Exploring effective vulnerability notifications,” in *25th USENIX Security Symposium (USENIX Security 16)*, 2016, pp. 1033–1050.
- [12] C. Rossow, “Amplification Hell: Revisiting Network Protocols for DDoS Abuse,” in *NDSS*, 2014.
- [13] Q. Lone, M. Korczyński, C. Gañán, and M. van Eeten, “Saving the internet: Explaining the adoption of source address validation by internet service providers,” in *Workshop on the Economics of Information Security*, 2020.
- [14] CAIDA, “The Spoofer Project,” 2020, <https://www.caida.org/projects/spoof/>.
- [15] T. M. Therneau and P. M. Grambsch, “The cox model,” pp. 39–77, 2000.
- [16] L. F. Müller, M. J. Luckie, B. Huffaker, kc claffy, and M. P. Barcellas, “Challenges in inferring spoofed traffic at IXPs,” in *ACM Conference on Emerging Networking Experiments And Technologies (CoNEXT)*, 2019, pp. 96–109.
- [17] F. Lichtblau, F. Streibelt, T. Krüger, P. Richter, and A. Feldmann, “Detection, Classification, and Analysis of Inter-domain Traffic with Spoofed Source IP Addresses,” in *Internet Measurement Conference*. ACM, 2017.
- [18] M. Korczyński, Y. Nosyk, Q. Lone, M. Skwarek, B. Jonglez, and A. Duda, “Don’t Forget to Lock the Front Door! Inferring the Deployment of Source Address Validation of Inbound Traffic,” in *Passive and Active Measurement Conference (PAM)*, 2020.
- [19] M. Korczyński, Y. Nosyk, Q. Lone, M. Skwarek, B. Jonglez, and A. Duda, “Inferring the deployment of inbound source address validation using dns resolvers,” in *Proceedings of the Applied Networking Research Workshop*, 2020, pp. 9–11.
- [20] R. Beverly, A. Berger, Y. Hyun, and k. claffy, “Understanding the Efficacy of Deployed Internet Source Address Validation Filtering,” in *Internet Measurement Conference*. ACM, 2009.
- [21] R. Beverly and S. Bauer, “The Spoofer Project: Inferring the Extent of Source Address Filtering on the Internet,” in *USENIX Steps to Reducing Unwanted Traffic on the Internet (SRUTI) Workshop*, Jul. 2005.
- [22] M. Kühner, T. Hüpperich, C. Rossow, and T. Holz, “Exit from Hell? Reducing the Impact of Amplification DDoS Attacks,” in *23th USENIX Security Symposium (USENIX Security 14)*, 2014.
- [23] Q. Lone, M. Luckie, M. Korczyński, and M. van Eeten, “Using Loops Observed in Traceroute to Infer the Ability to Spoof,” in *Passive and Active Measurement Conference*, 2017.
- [24] O. Cetin, C. Gañán, L. Altena, S. Tajalizadehkhooob, and M. van Eeten, “Tell me you fixed it: Evaluating vulnerability notifications via quarantine networks,” in *2019 IEEE European Symposium on Security and Privacy (EuroS&P)*. IEEE, 2019, pp. 326–339.
- [25] RDAP Client, 2020, <https://about.rdap.org/>.
- [26] B. Stock, G. Pellegrino, F. Li, M. Backes, and C. Rossow, “Didn’t you hear me?—towards more successful web vulnerability notifications,” 2018.
- [27] B. Stock, G. Pellegrino, C. Rossow, M. Johns, and M. Backes, “Hey, you have a problem: On the feasibility of large-scale web vulnerability notification,” in *25th USENIX Security Symposium (USENIX Security 16)*, 2016, pp. 1015–1032.
- [28] O. Cetin, M. Hanif Jhaveri, C. Gañán, M. van Eeten, and T. Moore, “Understanding the role of sender reputation in abuse reporting and cleanup,” *Journal of Cybersecurity*, vol. 2, no. 1, pp. 83–98, 2016.
- [29] O. Cetin, C. Ganan, M. Korczyński, and M. van Eeten, “Make notifications great again: learning how to notify in the age of large-scale vulnerability scanning,” in *Workshop on the Economics of Information Security (WEIS)*, 2017.
- [30] R. H. Thaler and C. R. Sunstein, “Libertarian paternalism,” *American economic review*, vol. 93, no. 2, pp. 175–179, 2003.
- [31] Thaler, Richard H. and Sunstein, Cass R., “Nudge improving decisions about health, wealth and happiness.” Penguin, 2021.
- [32] C. R. Sunstein, “Nudging: a very short guide,” in *The Handbook of Privacy Studies*. Amsterdam University Press, 2018, pp. 173–180.
- [33] E. Peer, S. Egelman, M. Harbach, N. Malkin, A. Mathur, and A. Frik, “Nudge me right: Personalizing online security nudges to people’s decision-making styles,” *Computers in Human Behavior*, vol. 109, p. 106347, 2020.
- [34] A. Frik, N. Malkin, M. Harbach, E. Peer, and S. Egelman, “A promise is a promise: The effect of commitment devices on computer security intentions,” in *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 2019, pp. 1–12.
- [35] A. Acquisti, I. Adjerid, R. Balebako, L. Brandimarte, L. F. Cranor, S. Komanduri, P. G. Leon, N. Sadeh, F. Schaub, M. Sleeper *et al.*, “Nudges for privacy and security: Understanding and assisting users’ choices online,” *ACM Computing Surveys (CSUR)*, vol. 50, no. 3, pp. 1–41, 2017.
- [36] M. Nagatsu, “Social nudges: their mechanisms and justification,” *Review of Philosophy and Psychology*, vol. 6, no. 3, pp. 481–494, 2015.
- [37] A. Brandon, P. J. Ferraro, J. A. List, R. D. Metcalfe, M. K. Price, and F. Rundhammer, “Do the effects of social nudges persist? theory and evidence from 38 natural field experiments,” National Bureau of Economic Research, Tech. Rep., 2017.
- [38] H. C. Kelman and V. L. Hamilton, *Crimes of obedience: Toward a social psychology of authority and responsibility*. Yale University Press, 1989.
- [39] R. B. Cialdini, “The psychology of persuasion,” *New York*, 1993.
- [40] Cialdini, Robert B, “The science of persuasion,” *Scientific American*, vol. 284, no. 2, pp. 76–81, 2001.
- [41] A. Falk and U. Fischbacher, “A theory of reciprocity,” *Games and economic behavior*, vol. 54, no. 2, pp. 293–315, 2006.
- [42] A. W. Gouldner, “The norm of reciprocity: A preliminary statement,” *American sociological review*, pp. 161–178, 1960.
- [43] J. Berg, J. Dickhaut, and K. McCabe, “Trust, reciprocity, and social history,” *Games and economic behavior*, vol. 10, no. 1, pp. 122–142, 1995.
- [44] E. Fehr and S. Gächter, “Fairness and retaliation: The economics of reciprocity,” *Journal of economic perspectives*, vol. 14, no. 3, pp. 159–181, 2000.
- [45] M. A. Nowak and K. Sigmund, “Evolution of indirect reciprocity,” *Nature*, vol. 437, no. 7063, pp. 1291–1298, 2005.
- [46] I. Seinen and A. Schram, “Social status and group norms: Indirect reciprocity in a repeated helping experiment,” *European economic review*, vol. 50, no. 3, pp. 581–602, 2006.
- [47] J. Mauch, “Spoofing ASNs,” 2013, <http://seclists.org/nanog/2013/Aug/132>.
- [48] MaxMind LLC, “MaxMind GeoIP,” 2020. [Online]. Available: <https://www.maxmind.com/en/geoip2-databases>
- [49] “PeeringDB,” 2020, <https://www.peeringdb.com>.
- [50] F. of Incident Response and S. Teams, 2020, <https://www.first.org/>.
- [51] Software Engineering Institute, 2020, <https://www.sei.cmu.edu/our-work/cybersecurity-center-development/national-csirts/>.
- [52] M. Kühner, T. Hüpperich, J. Bushart, C. Rossow, and T. Holz, “Going Wild: Large-Scale Classification of Open DNS Resolvers,” in *Internet Measurement Conference*. ACM, 2015.
- [53] A. Lodhi, N. Larson, A. Dhamdhere, C. Dovrolis, and K. Claffy, “Using peeringdb to understand the peering ecosystem,” *ACM SIGCOMM Computer Communication Review*, vol. 44, no. 2, pp. 20–27, 2014.

- [54] T. Böttger, F. Cuadrado, and S. Uhlig, "Looking for hypergiants in peeringdb," *ACM SIGCOMM Computer Communication Review*, vol. 48, no. 3, pp. 13–19, 2018.
- [55] A. Kühberger, "The framing of decisions: A new look at old problems," *Organizational Behavior and Human Decision Processes*, vol. 62, no. 2, pp. 230–240, 1995.
- [56] R. J. Donovan and G. Jalleh, "Positive versus negative framing of a hypothetical infant immunization: the influence of involvement," *Health Education & Behavior*, vol. 27, no. 1, pp. 82–95, 2000.
- [57] D. H. Rosenblatt, S. Bode, H. Dixon, C. Murawski, P. Summerell, A. Ng, and M. Wakefield, "Health warnings promote healthier dietary decision making: Effects of positive versus negative message framing and graphic versus text-based warnings," *Appetite*, vol. 127, pp. 280–288, 2018.
- [58] J. E. Grizzle, "A note on stratifying versus complete random assignment in clinical trials," *Controlled clinical trials*, vol. 3, no. 4, pp. 365–368, 1982.
- [59] J. Van Der Ham, "Ethics and internet measurements," in *2017 IEEE Security and Privacy Workshops (SPW)*. IEEE, 2017, pp. 247–251.
- [60] E. Kenneally and D. Dittrich, "The Menlo Report: Ethical Principles Guiding Information and Communication Technology Research," *SSRN Electronic Journal*, 2012.
- [61] "Understanding Relative Risk, Odds Ratio and Related Terms," 2020, <https://www.pitt.edu/~bertsch/risk.pdf>.
- [62] "Manrs for network operators," 2021, <https://www.manrs.org/isps/>.
- [63] "University of Oregon Route Views Project," 2020, <http://www.routeviews.org/routeviews/>.
- [64] TeleGeography, 2020, <https://www.telegeography.com/products/globalcomms/>.
- [65] CAIDA, "Macroscopic Internet Topology Data Kit (ITDK)," 2020, <https://www.caida.org/data/internet-topology-data-kit/>.
- [66] K. Keys, Y. Hyun, M. Luckie, and K. Claffy, "Internet-scale ipv4 alias resolution with midar," *IEEE/ACM Transactions on Networking*, vol. 21, no. 2, pp. 383–399, 2012.
- [67] A. Marder, M. Luckie, A. Dhamdhare, B. Huffaker, K. Claffy, and J. M. Smith, "Pushing the boundaries with bdrmapit: Mapping router ownership at internet scale," in *Proceedings of the Internet Measurement Conference 2018*, 2018, pp. 56–69.
- [68] X. Dimitropoulos, D. Krioukov, M. Fomenkov, B. Huffaker, Y. Hyun, G. Riley *et al.*, "AS relationships: Inference and Validation," *ACM SIGCOMM Computer Communication Review*, vol. 37, no. 1, pp. 29–40, 2007.
- [69] ITU, 2020, <http://www.itu.int/net4/ITU-D/idi/2017/index.html>.
- [70] "MANRS Implementation Guide," 2020, <https://www.manrs.org/isps/guide/antispoofing/>.
- [71] "RIPE 81," 2020, <https://ripe81.ripe.net/archives/video/420/>.
- [72] "RIPE Roundtable," 2020, <https://www.ripe.net/participate/meetings/roundtable/january-2017/presentations/security-and-the-ripe-community>.
- [73] "NANOG 2018," 2020, https://pc.nanog.org/static/published/meetings/NANOG2019/1838/20180921_Wittkop_Routing_Security_Ddos_v1.pdf.
- [74] "NANOG75," 2020, https://pc.nanog.org/static/published/meetings/NANOG75/1887/20190219_Compton_Ebgp_Flowspec_Peering_v1.pdf.
- [75] "NANOG 2018," 2020, <https://www.internetsociety.org/issues/manrs/>.
- [76] M. Vasek and T. Moore, "Do malware reports expedite cleanup? an experimental study," in *CSET*, 2012.
- [77] M. Vasek, M. Weeden, and T. Moore, "Measuring the impact of sharing abuse data with web hosting providers," in *Proceedings of the 2016 ACM on Workshop on Information Sharing and Collaborative Security*, 2016, pp. 71–80.
- [78] W. Soussi, M. Korczyński, S. Maroofi, and A. Duda, "Feasibility of large-scale vulnerability notifications after gdpr," in *2020 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*. IEEE, 2020, pp. 532–537.
- [79] E. Zeng, F. Li, E. Stark, A. P. Felt, and P. Tabriz, "Fixing https misconfigurations at scale: An experiment with security notifications," *Workshop on the Economics of Information Security*, 2019.
- [80] "RIPE IP Anti-Spoofing Task Force," 2021, <https://www.ripe.net/participate/ripe/tf/anti-spoofing>.
- [81] "NANOG75," 2020, https://pc.nanog.org/static/published/meetings/NANOG75/1956/20190219_Levy_Lightning_Talk_Dropping_v1.pdf.
- [82] "BGP hijacker booted off the Internet's backbone," 2020, https://www.theregister.com/2018/07/11/bgp_hijacker_booted_off_the_internets_backbone.
- [83] S. M. Diop, J. D. Ndibwile, D. Fall, S. Kashiwara, and Y. Kadobayashi, "To coerce or not to coerce? a quantitative investigation on cybersecurity and cybercrime legislations towards large-scale vulnerability notifications," in *2019 IEEE International Symposium on Software Reliability Engineering Workshops (ISSREW)*. IEEE, 2019, pp. 282–287.

APPENDIX A NOTIFICATION TEXT

A. Direct Notifications – Baseline

Subject : Possible IP spoofing from AS X

We are security researchers from Delft University of Technology. We have conducted a test to detect potential IP spoofing. **DETECTED ISSUE:** We have observed that your network may be allowing IP spoofing. You can check the test results at: [LINK]

WHAT TO DO: We encourage you to deploy Source Address Validation (BCP38) in your network today: <https://www.manrs.org/isps/guide/antispoofing/>

HOW TO VALIDATE: Please run the Spoofer tool to validate if BCP38 was implemented correctly: <https://www.caida.org/projects/spoofier/#software>

CONTACT: If you have any questions, concerns, issues, or comments, please send an email to infospoofing@tudelft.nl

B. Direct Notification – Social Nudge

Subject : Possible IP spoofing from AS X

We are security researchers from Delft University of Technology. We have conducted a test to detect potential IP spoofing. **DETECTED ISSUE:** We have observed that your network may be allowing IP spoofing. You can check the test results at: [LINK]

WHAT TO DO: We encourage you to deploy Source Address Validation (BCP38) in your network today: <https://www.manrs.org/isps/guide/antispoofing/>

Note that 75% of network operators in the world already deploy BCP38 in their networks. Deploy BCP38 in your network to become one of them.

HOW TO VALIDATE: Please run the Spoofer tool to validate if BCP38 was implemented correctly: <https://www.caida.org/projects/spoofier/#software>

CONTACT: If you have any questions, concerns, issues, or comments, please send an email to infospoofing@tudelft.nl

C. Direct Notification – Reciprocity

Subject : Possible IP spoofing from AS X

We are security researchers from Delft University of Technology. We have conducted a test to detect potential IP spoofing. **DETECTED ISSUE:** We have observed that your network may be allowing IP spoofing. You can check the test results at: [LINK]

WHAT TO DO: We encourage you to deploy Source Address Validation (BCP38) in your network today: <https://www.manrs.org/isps/guide/antispoofing/>

Note that your network is receiving fewer DDoS attacks because other networks have deployed BCP38. Return the favor - deploy BCP38 in your network to make the Internet more secure.

HOW TO VALIDATE: Please run the Spoofer tool to validate if BCP38 was implemented correctly: <https://www.caida.org/projects/spoofersoftware>

CONTACT: If you have any questions, concerns, issues, or comments, please send an email to infospoofing@tudelft.nl

D. CERT Notification – Baseline

Subject : Possible IP Spoofing from ASes in $\langle COUNTRY \rangle$
We are security researchers from Delft University of Technology. We have conducted a test to detect potential IP spoofing. We have observed that certain network operators in your country may be allowing IP spoofing. You can check the test results at: [LINK]

We encourage you to recommend those operators to deploy Source Address Validation (BCP38) in their network.

For your convenience, we tailored a draft of the notification for the network operators. This draft has been tested for clarity and comprehension and has been validated by the experts. We highly recommend you including this draft in your notification to the network operators.

DRAFT OF THE NOTIFICATION: Security researchers from Delft University of Technology have conducted a test to detect potential IP spoofing.

DETECTED ISSUE: They have observed that your network may be allowing IP spoofing. You can check the test results at: [LINK] (NOTE: Before sending out the notification, please insert the appropriate AS NUMBER)

WHAT TO DO: We encourage you to deploy Source Address Validation (BCP38) in your network today: <https://www.manrs.org/isps/guide/antispoofing/>

HOW TO VALIDATE: Please run the Spoofer tool to validate if BCP38 was implemented correctly: <https://spoofersoftware.caida.org/projects/spoofersoftware>

CONTACT: If you have any questions, concerns, issues, or comments, please send an email to infospoofing@tudelft.nl

E. CERT Notification–Social Nudge

Subject : Possible IP Spoofing from ASes in $\langle COUNTRY \rangle$
We are security researchers from Delft University of Technology. We have conducted a test to detect potential IP spoofing. We have observed that certain network operators in your country may be allowing IP spoofing. You can check the test results at: [LINK]

We encourage you to recommend those operators to deploy Source Address Validation (BCP38) in their network.

For your convenience, we tailored a draft of the notification for the network operators. This draft has been tested for clarity and comprehension and has been validated by the experts. We highly recommend you including this draft in your notification to the network operators.

DRAFT OF THE NOTIFICATION: Security researchers from Delft University of Technology have conducted a test to detect potential IP spoofing.

DETECTED ISSUE: They have observed that your network may be allowing IP spoofing. You can check the test results

at: [LINK] (NOTE: Before sending out the notification, please insert the appropriate AS NUMBER)

WHAT TO DO: We encourage you to deploy Source Address Validation (BCP38) in your network today: <https://www.manrs.org/isps/guide/antispoofing/>

Note that 75% of network operators in the world already deploy BCP38 in their networks. Deploy BCP38 in your network to become one of them.

HOW TO VALIDATE: Please run the Spoofer tool to validate if BCP38 was implemented correctly: <https://www.caida.org/projects/spoofersoftware>

CONTACT: If you have any questions, concerns, issues, or comments, please send an email to infospoofing@tudelft.nl.

F. CERT Notification–Reciprocity

Subject : Possible IP Spoofing from ASes in $\langle COUNTRY \rangle$
We are security researchers from Delft University of Technology. We have conducted a test to detect potential IP spoofing. We have observed that certain network operators in your country may be allowing IP spoofing. You can check the test results at: [LINK]

We encourage you to recommend those operators to deploy Source Address Validation (BCP38) in their network.

For your convenience, we tailored a draft of the notification for the network operators. This draft has been tested for clarity and comprehension and has been validated by the experts. We highly recommend you including this draft in your notification to the network operators.

DRAFT OF THE NOTIFICATION: Security researchers from Delft University of Technology have conducted a test to detect potential IP spoofing.

DETECTED ISSUE: They have observed that your network may be allowing IP spoofing. You can check the test results at: [LINK] (NOTE: Before sending out the notification, please insert the appropriate AS NUMBER)

WHAT TO DO: We encourage you to deploy Source Address Validation (BCP38) in your network today: <https://www.manrs.org/isps/guide/antispoofing/>

Note that your network is receiving fewer DDoS attacks because other networks have deployed BCP38. Return the favor - deploy BCP38 in your network to make the Internet more secure.

HOW TO VALIDATE: Please run the Spoofer tool to validate if BCP38 was implemented correctly: <https://www.caida.org/projects/spoofersoftware>

CONTACT: If you have any questions, concerns, issues, or comments, please send an email to infospoofing@tudelft.nl.

G. NOG Notification

CAIDA's source address validation measurement project (<https://spoofersoftware.caida.org>) is automatically generating monthly reports of ASes originating prefixes in BGP for systems from which we received packets with a spoofed source address.

We are publishing these reports to network and security operations lists in order to ensure this information reaches

operational contacts in these ASes. This report summarises tests conducted within *<COUNTRY>*.
Inferred improvements during *<DATE>*:

ASN	Name	Fixed by
ASN	ASN NAME	DATE

Further information for the inferred remediation is available at: <https://spoofer.caida.org/remedy.php>
Source Address Validation issues inferred using Spoofer tool during *<DATE>* :

ASN	Name	First-Spoofed	Last-Spoofed
ASN	ASN NAME	DATE	DATE

Further information for these tests where we received spoofed packets using spoofer is available at: https://spoofer.caida.org/recent_tests.php?country_include=ccc,ccc&no_block=1
Source Address Validation issues inferred using misconfigured open resolvers during *<DATE>*:

ASN	Name	First-Spoofed	Last-Spoofed
ASN	ASN NAME	DATE	DATE

Further information for these tests where we received spoofed packets using open resolver is available at:[LINK]
Please send any feedback or suggestions to [spoofer-info at caida.org](mailto:spoofer-info@caida.org)

APPENDIX B QUESTIONNAIRE

Q1: In your opinion, does your network have any of the following security issues? Choose all that apply.

- 1) Susceptible to Route/Prefix Hijack
- 2) Does not prevent IP spoofing
- 3) Susceptible to DDoS
- 4) None of the above
- 5) I'm not sure

Q2: How did you discover the issue with IP spoofing? Choose all that apply.

- 1) I ran a Spoofer test
- 2) I received a notification from NOG (Network Operator Group)
- 3) I received a notification from CERT (Computer Emergency Response Team)
- 4) I received a notification from security researchers
- 5) Other (please specify)

Q3: Are you the person responsible for the implementation of Source Address Validation (SAV), which is also referred to as BCP38?

- 1) Yes
- 2) No
- 3) I'm not sure
- 4) I don't know what SAV means

Q4: Have you escalated the issue with IP spoofing to the person/team responsible for SAV implementation?

- 1) Yes
- 2) No
- 3) I'm not sure

Q5: Have you implemented SAV in your network?

- 1) Yes, on the entire network
- 2) Yes, but only in the segment of our network
- 3) No, we haven't implemented SAV in our network at all
- 4) I'm not sure

Q6: What kind of filtering of origin IPs do you perform? Choose all that apply.

- 1) Filter private address space (RFC 1918)
- 2) Perform SAV on customer facing interfaces
- 3) Perform SAV on stub AS
- 4) Other (please specify)

Q7: Why didn't you implement SAV in your network? Choose all that apply.

- 1) I lack technical knowledge to implement SAV
- 2) I am concerned that SAV implementation may cause network downtime/performance
- 3) I don't have time to implement SAV at the moment
- 4) I don't think IP spoofing is an important issue
- 5) I don't think DDoS (Distributed Denial of Service Attack) is an important issue
- 6) I don't think SAV is effective in addressing IP spoofing issues
- 7) We are running a non-stub network
- 8) We are running a multi-homed network
- 9) Other (please specify)

Q8: Are you planning to implement SAV in your network?

- 1) Yes
- 2) No
- 3) I'm not sure

Q9: MANRS provides the following guidelines for implementing SAV: <https://www.manrs.org/isps/guide/antispoofing/>. Please review the guidelines and tell us your opinion: Do you think the MANRS guidelines provide sufficient information on how to implement SAV in your network?

- 1) Yes
- 2) No
- 3) I'm not sure

Q10:What information, necessary for implementing SAV, is missing in MANRS guidelines? Please, provide as much details as you can.

APPENDIX C SCREEN SHOT OF WEBSITE

Below is an example for website linked to the notification to AS137612

This page contains evidence that a network may not have deployed Source Address Validation (SAV) to block packets with source addresses that are invalid given the attachment point.

Our method to detect a possible lack of SAV is based on querying Open Resolvers. When we queried the Open Resolver resolver IP address listed, we received a response to the query from the listed Recursive Resolver IP that maps to a different ASN, which is likely not a valid source address for that network attachment point.

Each row contains a link to a report with further details of how we observed a possible lack of SAV for that Open Resolver.

Id	Timestamp (UTC)	Open Resolver IP	Open Resolver ASN	Country	Recursive Resolver IP	Recursive Resolver ASN
5566796	2020-12-05 03:48:05	103.117.38.214	137612 (CDCN-AS-IN)	ind	8.8.8.8	15169 (GOOGLE)
5566784	2020-12-05 03:47:42	103.117.39.236	137612 (CDCN-AS-IN)	ind	8.8.8.8	15169 (GOOGLE)
5566359	2020-12-05 03:28:58	103.117.38.227	137612 (CDCN-AS-IN)	ind	8.8.8.8	15169 (GOOGLE)

Fig. 9. Main page with individual reports per IP address for AS137612

Description:

This page describes the outcome of probing the open resolver with IP address 103.117.38.214 in AS 137612. When we sent that open resolver a DNS query with a domain name under our authoritative control, we received a response from 8.8.8.8. We present two possible explanations for the behavior that indicate the network hosting the open resolver has not deploying source address validation.

You may be able to reproduce these results using the dig command, as follows:

```
$ dig www.example.net @103.117.38.214
;; reply from unexpected source: 8.8.8.8#53, expected 103.117.38.214
```

Summary:

Timestamp: 2020-12-05 03:48:05
 Open Resolver IP: 103.117.38.214
 Open Resolver ASN: 137612
 Open Resolver Country: ind
 Recursive Resolver IP: 8.8.8.8
 Recursive Resolver ASN: 15169

Fig. 10. Details about our methodology and steps to reproduce the results

Second case: Open Resolver forwards response to our Vantage Point without rewriting Source Address

In the second case, the open resolver forwards the query and receives the response from its configured recursive resolver. However, the open resolver then forwards the response back to our vantage point, without rewriting the source address of the reply from the recursive resolver. Because the IP address of the recursive resolver is not a valid source IP address at the network attachment point hosting the open resolver, the response should be filtered if the network hosting the open resolver has configured source address validation.

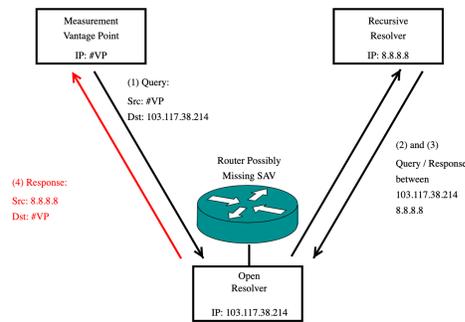


Fig. 12. Explanation of the second case with dynamic IP addresses for the figures

First case: Open Resolver forwards query to Recursive Resolver without rewriting Source Address

In the first case, (1) the open resolver forwards the query directly to its configured recursive resolver without rewriting the source IP address from the Vantage Point (VP) we sent the query from, and (2) we subsequently receive the response to our query directly from the recursive resolver. Because the source IP address is not valid at the network attachment point hosting the open resolver, the query should be filtered if the network hosting the open resolver has configured source address validation.

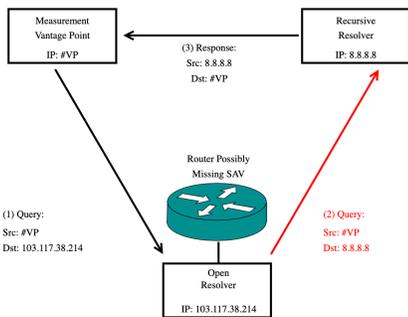


Fig. 11. Explanation of the first case with dynamic IP addresses for the figures