

Network Hygiene, Incentives, and Regulation: Deployment of Source Address Validation in the Internet

Matthew Luckie
University of Waikato
mjl@wand.net.nz

Robert Beverly
Naval Postgraduate School
rbeverly@nps.edu

Ryan Koga
CAIDA, UC San Diego
rkoga@caida.org

Ken Keys
CAIDA, UC San Diego
kkeys@caida.org

Joshua A. Kroll
Naval Postgraduate School
jkroll@jkroll.com

k claffy
CAIDA, UC San Diego
kc@caida.org

ABSTRACT

The Spoofer project has collected data on the deployment and characteristics of IP source address validation on the Internet since 2005. Data from the project comes from participants who install an active probing client that runs in the background. The client automatically runs tests both periodically and when it detects a new network attachment point. We analyze the rich dataset of Spoofer tests in multiple dimensions: across time, networks, autonomous systems, countries, and by Internet protocol version. In our data for the year ending August 2019, at least a quarter of tested ASes did not filter packets with spoofed source addresses leaving their networks. We show that routers performing Network Address Translation do not always filter spoofed packets, as 6.4% of IPv4/24 tested in the year ending August 2019 did not filter. Worse, at least two thirds of tested ASes did not filter packets entering their networks with source addresses claiming to be from within their network that arrived from outside their network. We explore several approaches to encouraging remediation and the challenges of evaluating their impact. While we have been able to remediate 352 IPv4/24, we have found an order of magnitude more IPv4/24 that remains unremediated, despite myriad remediation strategies, with 21% unremediated for more than six months. Our analysis provides the most complete and confident picture of the Internet's susceptibility to date of this long-standing vulnerability. Although there is no simple solution to address the remaining long-tail of unremediated networks, we conclude with a discussion of possible non-technical interventions, and demonstrate how the platform can support evaluation of the impact of such interventions over time.

CCS CONCEPTS

• **Networks** → **Network security**.

KEYWORDS

IP spoofing; remediation

ACM Reference Format:

Matthew Luckie, Robert Beverly, Ryan Koga, Ken Keys, Joshua A. Kroll, and k claffy. 2019. Network Hygiene, Incentives, and Regulation: Deployment of Source Address Validation in the Internet. In *2019 ACM SIGSAC Conference on Computer and Communications Security (CCS '19)*, November 11–15, 2019, London, United Kingdom. ACM, New York, NY, USA, 16 pages. <https://doi.org/10.1145/3319535.3354232>

1 INTRODUCTION

IP source address spoofing is the process of generating IP packets with arbitrary source addresses, i.e., addresses other than those assigned to a host based on its network interface attachment point. Hosts can trivially generate spoofed-source IP packets. Malicious actors exploit this spoofing ability to mount a wide variety of attacks, e.g., volumetric denial-of-service [26] (DoS), resource exhaustion [17], policy evasion [39], and cache poisoning [53] to name just a few. In April 2019, IP addresses of large U.S. bank websites were spoofed by an attacker that used them to perform suspicious scanning [36] so that the addresses appeared on blacklists. This creative use of spoofing caused security products to block the bank's addresses, such that people using those security products, even unknowingly, were unable to interact with their banks.

Highly distributed ownership of network infrastructure makes it operationally difficult to block or trace back attacks using spoofed addresses to their true source. Therefore, best common practice for nearly 20 years has enjoined operators to verify the source addresses of traffic leaving their networks. Commonly referred to as "Source Address Validation" (SAV) or Best Current Practice (BCP) 38 [19], this prophylactic only prevents a provider who deploys SAV from originating spoofed-source traffic; it does not protect the provider from receiving spoofed traffic or being the victim of an attack. Unfortunately, continual incidences of spoofing demonstrates that SAV is not ubiquitously deployed. Spoofing continues to serve as a primary attack vector for large-scale DoS attacks [3, 27], and these attacks continue to increase in prevalence [24] and intensity; in 2018 GitHub experienced attacks of 1.35Tbps [26].

In this work, we report on long-term efforts and results of the Spoofer project. The Spoofer project is an effort to crowd-source measurements of the ability to spoof from many points in the network, and thus better understand the Internet's susceptibility to spoofed-source attacks. The data from the Spoofer project comes from volunteers who run the Spoofer client, which sends and receives a variety of spoofed packets. On the basis of which packets sent by the client are received by servers maintained by the Spoofer project, and which packets sent by the servers are received by the

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CCS'19, November 11–15, 2019, London, United Kingdom

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6747-9/19/11...\$15.00

<https://doi.org/10.1145/3319535.3354232>

client, the system infers the granularity and types of any SAV on paths involving the client.

The Spoofer project’s primary goal is to serve as an independent auditor and long-term record of Internet-wide SAV deployment. Toward this goal, we have continually improved the project by: i) removing barriers and incentivizing spoof testing; ii) making changes to the system to gather more tests from more locations; iii) adding tests that deepen our understanding of SAV deployment; and iv) attempting to incentivize SAV deployment. The data we have amassed represents the most comprehensive picture of SAV deployment on the Internet currently available. On the basis of this data, we report on the following five contributions:

(1) Three years of longitudinal Spoofer measurements collected by an automated client. In addition to reporting on previously uninvestigated aspects of SAV, e.g., IPv6 spoofing ability, spoofing through Network Address Translation (NAT) devices, and filtering inbound into a destination network, we perform a macro-level analysis of the Internet’s resistance to spoofing along multiple dimensions. Despite obtaining significantly more tests (both across time and topology), we find that the prevalence of SAV filtering has not measurably increased in the past decade. (§4)

(2) Quantitative assessment of the representativeness of the data. Crowd-sourced measurements present inherent challenges to survey data because participants impart bias. While our system design (§3) removes barriers to testing, and permits continual gathering of tests, we observe a decidedly non-uniform test coverage across networks and geographic regions. We therefore examine the extent to which the daemonized client successfully gathers longitudinal data. We build a simple model to predict spoofability based on previously observed measurements, and use it as a proxy for the quality of the data we have gathered. By showing that our model yields accurate predictions, we gain confidence in the degree to which our results have predictive power and reflect the larger Internet. (§5)

(3) A comprehensive understanding of the relationship between NAT as SAV, and the implications of an IPv6 Internet without NAT. Challenging a commonly held assumption that NATs prevent spoofing, we show that clients in 6.4% of IPv4 prefixes tested in the year ending August 2019 were able to send packets with spoofed source addresses from behind a NAT, and these packets were not filtered by their ISP. Not only do NATs not prevent spoofing, but the deployment of IPv6 presents new opportunities for attackers: many inexpensive, vulnerable IoT devices connected without NATs, capable of spoofing addresses from a much larger address space. We characterize SAV in both the context of NAT and IPv6 to dispel misconceptions about their role in abuse. (§6)

(4) Analysis of concerted remediation efforts, including publishing (“name-and-shame”) lists of providers with missing or misconfigured SAV. Between February 2016 and December 2018, we sent 1,877 private email notifications to networks that failed the SAV test. Beginning April 2018, we sent geographically-scoped public emails to regional network operator group mailing lists. After we stopped sending private notifications, the rate of remediation did not drop, leading us to believe that the private notifications had no measurable impact on remediation. (§7)

(5) Discussion of practical steps to increase global SAV deployment. Our work demonstrates the difficulty of incentivizing

providers to deploy SAV. However, we find several areas of “low hanging fruit” that are incentive-compatible and would have significant impact if adopted. Specifically, we show that operators can protect their own networks by filtering spoofed packets claiming to be from within their network when they arrive from outside of their network, and we highlight the set of Autonomous Systems (ASes) that are conducive to their provider’s use of filtering using Unicast Reverse Path Forwarding (uRPF), with no collateral damage. We include possible non-technical interventions, and demonstrate how the platform can support evaluation of the impact of such interventions over time. We argue that the only likely way to cover this long-tail of remediation is for equipment manufacturers to enable SAV by default. (§8)

2 RELATED WORK

On the prevalence of spoofed-source DoS attacks. In 2000 and 2001, spoofed-source attacks were prevalent enough for researchers to propose methods to trace back the source of spoofed packets [47, 49]; none have seen deployment due to operational and coordination costs. Recently, Jonker *et al.* analyzed data sets that covered a two-year period (March 2015 to February 2017) to infer 20 million denial of service attacks targeting 2.2 million /24 IPv4 blocks, more than one-third of those estimated to be active on the Internet [24]. A 2017 survey of 84 operators [30] confirmed the lack of resources (both knowledge and time) required to accurately maintain SAV filtering. The more fundamental issue is misaligned incentives: namely, SAV benefits other people’s networks (and their customers), not the network that has to deploy it (or its customers).

On promotion of SAV deployment. Many academic research efforts have described techniques to promote deployment of SAV [15, 31, 32, 62]. In 2014, the Internet Society began to foster grassroots community support to launch the global MANRS initiative – Mutually Agreed Norms for Routing Security [22], which included a public commitment to deploy source address validation, among other routing security best practices. In §8 we show that, in our data, their members are no more likely than the general population to deploy SAV. In 2016, the U.S. National Institute for Standards and Technology (NIST) provided a technical evaluation of the performance impact of deploying various types of reverse path filtering in commercial routers [38], and in 2018 provided deployment guidance [50].

On crowd-sourced measurement of SAV deployment. In 2005, Beverly *et al.* developed a client-server technique to allow users to test networks to which they are currently attached [5], and operationalized a platform to track trends from February 2005 to April 2009 [6]. This system required a user to download and execute the client software once per measurement, limiting coverage. In 2017, Lone *et al.* used five paid crowd-sourcing platforms to collect SAV measurements over a six-week period [33]. They paid platform fees of \approx 2,000 Euros to have workers execute 1519 Spoofer measurements from 91 countries and 784 unique ASes, 342 of which the Spoofer project had not measured. They reported that the observed spoofability in these measurements was similar to the volunteer-based Spoofer system.

On inference of spoofing from other data sources. To overcome the requirement for a vantage point in every network, over the

last few years researchers have investigated creative opportunistic techniques to infer lack of SAV in other macroscopic Internet data sets. In 2017, Lone *et al.* reported a technique to infer evidence of spoofed traffic in massive traceroute archives, based on the knowledge that an edge network should not appear to provide transit in a traceroute path [34]. This method is also limited by what appears in the traceroute archives, as well as by the inconsistent addressing conventions used in traceroute implementations.

Also in 2017, Lichtblau *et al.* used a large Internet Exchange Point (IXP) as a vantage point for inferring which networks had not deployed SAV [30]. They compared source IP addresses of packets an IXP member sent across the IXP fabric with the range of source addresses expected based on routes observed using the Border Gateway Protocol (BGP). This approach faced challenges of accurately inferring ranges of source addresses expected for a given AS, and obtaining cooperation of the IXPs to access traffic data. The authors did not use the data for remediation.

Effectiveness of remediation attempts. Several studies have shown the difficulty in effecting remediation of vulnerabilities via notification, even for vulnerabilities that pose risks to the notified networks themselves, as opposed to the rest of the Internet. Stock *et al.* found a marginal improvement in remediation rate due to notification of 44,000 vulnerable web sites; a key obstacle was getting the notification to the right person in the organization [52]. Li *et al.* found a similarly daunting result, that the most effective approach was to notify contacts registered in WHOIS [14] via a message with detailed information about a vulnerability [28]. Such notifications had a statistically significant impact on improving remediation (11% more contacts remediated than in the control group), but only a minority took any remediative action, often only partial, and repeat notifications had no effect. Reporting vulnerabilities through Computer Emergency Response Team (CERT) organizations appeared to be of limited utility. Hastings *et al.* found even more disappointing results in their study of the response to disclosure of an RSA private key compromise for 0.5% of HTTPS-enabled Internet hosts [20]. In a different study on remediation of hijacked websites, Li *et al.* reported that browser interstitials and search engine warnings correlated with faster remediation compared to private notification via WHOIS contact alone [29].

3 SYSTEM ARCHITECTURE

The Spoofer system consists of client software, a set of vantage points for receiving spoofed packets sent by the client, and a server for coordinating measurements and sending spoofed packets to the clients, illustrated in figure 1. Measurements collected prior to 2016 were sparse across both time and space (topology) [6]. Because users had to manually run a program requiring root access, many prefixes and ASes had only a single test run from them. Since August 2015, we undertook development efforts to reduce or remove barriers to operating the measurement client and obtaining measurement data, and outreach efforts to promote both testing and deployment of SAV. Table 1 summarizes the major improvements we made, and some of the outreach to network operators at industry meetings.

First, the client now has a digitally signed install package available for all major desktop operating systems (Windows, MacOS, and Linux) as well as OpenWRT, and is also provided as open source.




Figure 1: Overview of Spoofer project architecture, consisting of client software, a server to coordinate measurements, and a series of receivers to collect packets sent by clients.

May 2015	Project re-launched.
Feb 2016	Client allows user to share test results publicly.
May 2016	New client released, with GUI and automated probing via daemon.
May 2016	Begin sending private notifications.
Oct 2016	Present at North American Network Operators' Group (NANOG) meeting.
Dec 2016	Probing speed improvements in client.
Mar 2017	Begin testing SAV inbound to client. Begin testing granularity of IPv6 SAV filters.
Oct 2017	Begin testing location of IPv6 SAV filters.
Apr 2018	Begin monthly geo-scoped emails to network operator groups.
Jun 2018	Release OpenWRT client.
Sep 2018	Deploy new NAT testing mode.
Dec 2018	Halt private notifications.

Table 1: Timeline of Spoofer project since re-launch.

The client includes a GUI to present results in English, communicates extensively and securely to our server using protobuf and TLS, and has the ability to automatically update itself. Ideally the project would provide a Spoofer App for Android or iOS, but neither platform provides the ability to construct packets with spoofed source addresses unless the device is jailbroken.

Second, instead of requiring a user to manually initiate tests, the client is now a daemon that runs tests automatically, periodically testing at most weekly by default, and also whenever the client attaches to a network it has not tested previously. The daemon increases not only the spatial coverage of the data, but also the temporal coverage, allowing richer and more confident inferences (§5). The client performs a handshake to receive a list of work items to do, such as send spoofed packets to a set of cooperating addresses, conduct traceroutes towards those addresses, and listen for spoofed packets with specific source IP addresses that the server sends to the client to test filtering of packets inbound to the client network.




Figure 2: Growth in coverage of IP prefixes and ASes measured by the Spoofer project over time, from the beginning of the project in 2005 until August 2019. The gaps are due to hardware failures. The vertical line at May 2016 indicates when the new daemonized client was released.

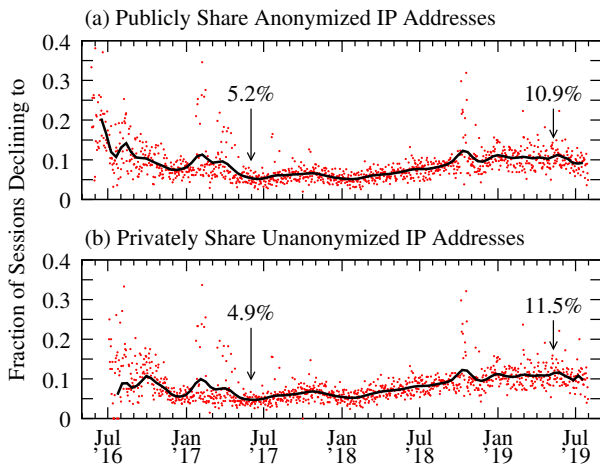


Figure 3: Fraction of Spoofer project tests with the daemonized client per day with sharing restrictions, over time. The daemonized client was released in May 2016. The Bézier curve shows the trend that the percentage of private tests grew from 5.2% of tests in June 2017, to 10.9% in May 2019.

Third, we expanded the testing capabilities of the client software. The handshake allows the server to determine if the client is operating from behind a Network Address Translation (NAT) router, and our system includes novel tests to evaluate SAV limitations in NAT implementations (§6.1). The client software now tests both IPv4 and IPv6 on all platforms, including SAV granularity and location in the network (§6.2). Our system also tests filtering of spoofed packets inbound to the client’s network, using packets from the Spoofer server with a source address from the client’s network – 192.0.2.0/24 in figure 1. The presence of a NAT prevents our testing of the client’s ability to receive spoofed packets because the NAT device cannot know to forward these packets to the client system.

Finally, to promote SAV deployment, the project now publicly shares test results with anonymized IP addresses on the project website, and privately shares unanonymized IP addresses with

Country	IPv4 and IPv6 ASes	
	Tested	Spoofable
United States	1316	27.6%
Brazil	494	55.1%
Netherlands	255	29.4%
Great Britain	217	29.5%
Russia	202	19.3%
Canada	171	29.2%
Germany	167	28.1%
India	155	27.1%
Australia	148	25.0%
Bangladesh	131	23.7%
Other:	2409	28.8%
Total:	5178	31.5%

Table 2: Number of ASes with prefixes geolocated to a given country between May 2016 and August 2019 in our data. We show the percentage of ASes geolocated to each country, and the percentage of those ASes that were spoofable.

operators who require unanonymized IP addresses to identify specific equipment without SAV deployed, provided the user does not opt-out of doing so. The Spoofer client prompts the user for their sharing preferences when it is first launched. For public reporting, we anonymize IP addresses by concealing at least the last eight bits of the address – i.e., for IPv4 addresses we publicly provide the first 24 bits of the address. For IPv6 addresses, we publicly provide the first 40 bits, as ISP operators may use DHCP prefix delegation to delegate prefixes between 48 and 64 bits to individual subscribers [40]. Beginning April 2018, we sent geographically-scoped public emails to regional network operator group mailing lists to encourage remediation when we had tests from ASes in the geographic region showing SAV had not been deployed in at least some networks in the region. To mitigate privacy concerns that might hinder adoption and use of the Spoofer client, client reports to the Spoofer system do not include any unique identifier.

The project’s efforts to lower barriers to running the test have had a measurable impact on coverage of Spoofer measurements (figure 2). Prior to the release of the probing daemon in May 2016, the project collected approximately 23k successful test sessions across 11 years, an average of 174 sessions per month. In the following 39 months ending August 2019, the project collected more than 414k sessions, averaging 10,639 sessions per month, i.e., two orders of magnitude more sessions per month than before the probing daemon was released, resulting in the project acquiring 17.75 times more data in 29.5% of the time. Measurement coverage of the global Internet topology has similarly increased, from 3410 IPv4 ASes to 6938 – 10.6% of the globally routed ASes as of August 2019 – with 3528 ASes newly observed since May 2016. IPv6 coverage has increased from 211 ASes to 980 – 6.3% of globally routed IPv6 ASes as of August 2019 – with 769 ASes newly observed since May 2016. Figure 3 shows that 5-10% of users opt-out of publicly sharing their test results, or privately sharing IP addresses with operators. The growth in private tests beginning in April 2018 is correlated with when the Spoofer project began sending geographically-scoped public emails to regional network operator group mailing lists, suggesting these public emails prompted operators to use the tool themselves, but that they did not want their results to be public.

The Spoofer system annotates each test by geolocating the client’s IP address at collection time using NetAcuity Edge. Table 2 shows the top 10 countries by number of ASes represented in our data, where each AS had an IPv4 or IPv6 prefix geolocated to the country that a client tested since the probing daemon was released in May 2016. Overall, we have tested at least one prefix in 5,178 ASes, with 31.5% of these ASes having at least one prefix that was spoofable. 34.9% of ASes were geolocated to one of two countries: the U.S., with 25.4% of the tested ASes, and Brazil with 9.5%. More than half of the ASes tested within Brazil did not block spoofed packets in at least one test, making Brazil an outlier in our data. We discuss remediation activity within the U.S. and Brazil in §7.

4 OVERVIEW

Figure 4 provides an overview of our findings on SAV deployment, for packets both outbound from and inbound to the measured network, for the year ending 1 August 2019. We present deployment statistics by IPv4/24 and IPv6/40 prefixes, and by AS. An AS with partial deployment originates some prefixes from which the Spoofer system did not receive spoofed packets, and originates other prefixes from which the Spoofer system did.

This data indicates that while most networks block packets with spoofed source IP addresses, deployment is far from universal. In particular, 25.2% and 32.1% of tested IPv4 and IPv6 ASes, respectively, had at least one prefix where operators had not deployed SAV to block outbound spoofed packets to the Internet (first and last bars in figure 4b). The comparatively small fraction of prefixes from which we received spoofed IPv4 packets (figure 4a) is primarily due to the presence of Network Address Translation (NAT) routers that rewrite spoofed packets with the router’s publicly routable address. When a NAT router was not present, 14.9% of IPv4 prefixes had no filtering; filtering was better deployed at prefix-level granularity in IPv6, with 12.3% of tested prefixes not filtering.




Figure 4: Summary statistics for Spoofer Project for the year ending August 2019. Percent values to right of barplot indicate the fraction of networks with problematic SAV deployments, either no filtering or partial filtering of spoofed packets (red plus orange segments).

The lower panels of figure 4 (c and d) summarize the observed state of filtering of packets *inbound* to the client’s network, claiming to be from within the same subnet as the client. This test sends a packet to the client with a source address that is inside the same subnet as the client. For IPv4 we toggle the last (31st) bit of the address, and for IPv6 we toggle the 120th bit. Surprisingly, inbound filtering of spoofed packets is even less deployed than outbound filtering, despite these packets being a threat to the receiving network. Deploying this type of filtering is incentive-compatible because internal hosts are often more trusted than external hosts [46]. In our data, 67.0% and 74.2% of IPv4 and IPv6 ASes, respectively, had at least one prefix where they were not filtering inbound packets.

We compare the IPv4 tests we received for the year ending August 2019 where the client was not behind a NAT, to the same class of tests 1 May 2006 to 7 May 2006, during which the project received 1057 tests, triggered by an article on Slashdot encouraging readers to run the original Spoofer test [61] (lower left panel of figure 2). For the first week of May 2006, 18.3% of IPv4/24 prefixes and 20.4% of ASes tested did not block spoofed-source packets from leaving their networks. Figure 4 shows that for the year ending August 2019, 14.9% of IPv4/24 prefixes and 30.5% of ASes tested did not block spoofed-source packets, implying that SAV deployment has not generally improved.

We reinforce this result by considering the influence of a prefix’s age – defined as how long it has been observable in the global BGP routing table – on the probability that it performs source address validation. One hypothesis is that older networks are more




Figure 5: For Spoofer tests collected between May 2016 and August 2019, we found no correlation between prefix age (measured as first appearance in the global BGP routing table) and its probability of SAV deployment (as measured by a client in that prefix performing a test).

mature and hence have better hygiene. However, newer networks may have newer equipment, less configuration inertia, and better awareness of network security. The increasing prevalence of address transfers post-IPv4-exhaustion adds additional noise to this already speculative exercise.

We generate a radix trie of prefix ages by iterating over monthly snapshots of Routeviews [1] and RIPE RIS from January 2000 to August 2019. The meta-data in the radix trie for each prefix is the date it first appeared in our snapshots. Note that prefix aggregation and deaggregation over time result in new ages within the radix trie – i.e., a deaggregation is a new BGP policy and hence resets the age. Let $t_0(p)$ be the time at which the IPv4 prefix p first appeared in the global BGP routing table. For a client with IPv4 address c who performs a Spoofer test at time t , we find the longest matching prefix $p' = LPM(c, t)$ within the radix trie at time t . The prefix age is then relative to the time of the test:

$$age(c, t) = t - t_0(LPM(c, t))$$

For each Spoofer client using a routed source address, we determined the age of the corresponding client prefix. We binned results into months and determined the relative fraction of tests within each age bin that could successfully spoof the routed IPv4 address. Figure 5 shows the relationship between the fraction of tests where a client could spoof and the age of the prefix to which the client belonged. Using Pearson’s correlation, we found no meaningful relationship between age and spoofability.

5 EPISTEMOLOGICAL CHALLENGES WITH CROWD-SOURCING MEASUREMENT

A crowd-sourced approach to assessing the deployment of SAV is confounded by multiple factors. First, the opt-in nature of the measurement can induce sample bias: people interested in security issues are more likely to provide a measurement. Second, and related, the measurement results we receive are distributed non-uniformly across time and networks. Third, the Internet itself is changing, e.g., due to equipment upgrades, configuration changes, and policy changes. Fourth, the results of a single test may not be indicative of the larger network, prefix, or autonomous system from which the client executes its test.




Figure 6: Even with periodic weekly probing by the client, only 8% of IPv4 /24 prefixes and 23% of ASes reported measurements in 20 or more weeks in data we collected between May 2016 and August 2019.

This sparsity complicates inferences of remediation. For instance, if two samples from an AS are significantly separated in time and from different prefixes, and the first sample indicates the tested network permits spoofing, while the second indicates filtering, there could be several different explanations for the change. The tested network might have deployed SAV across their network, fixing the first prefix; or the first tested prefix might still permit spoofing.

5.1 Effect of daemonizing Spoofer client

We attempted to mitigate sampling concerns by daemonizing the Spoofer client, i.e., running the client in the background after installation, executing measurements any time the client detected a new attached network, and repeating tests weekly on previously seen networks. In addition to obtaining more test samples, daemonizing the Spoofer client allows the system to automatically gather longitudinal samples of the same network – data that is useful to characterize the evolution of filtering policies.

Because we do not use a client identifier (§3) it is not possible to determine whether a specific client continues to probe after its initial installation. Instead, we must estimate the ability of the periodic daemon to gather longitudinal data for a given /24 prefix, longest matching BGP prefix, or AS. Figure 6 displays the cumulative fraction of these as a function of the number of active weeks. We define an active week to be any week for which there is a test report from a client with the daemon functionality.

Unfortunately, even with the daemon, 44% of the IPv4 /24 prefixes and 32% of the ASes reported measurements in only a single week between May 2016 and August 2019. This may be due to either users who run the tool once and then uninstall it after obtaining their test result, or infrequent or one-off tests performed by clients while they are roaming, e.g., mobile hosts that attach infrequently to a hotel or airport network. Although unexpected, these results are also consistent with feedback from operators that told us they used the client on a short-term basis while they were testing SAV configuration and then uninstalled the software from their laptop.

While our coverage is longitudinally sparse for many networks, figure 6 exhibits long tails. In particular, for large providers, we have tests in every week. For example, we obtained test results from clients in AT&T, Comcast, Verizon, Deutsche Telekom, Charter, and Cox in nearly all of the weeks since the release of the daemonized client – May 2016 to August 2019.


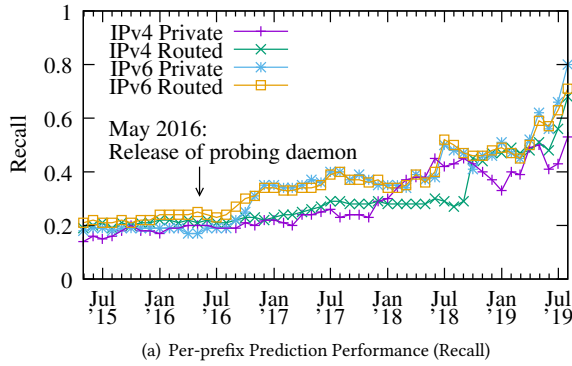


Figure 7: To assess the predictive power of our data-driven model of Internet spoofability, we trained a model based on client tests up to time t , and then used the model to predict test outcomes after t . The improvement in precision and recall after the release of the probing daemon in May 2016 gives us some confidence that our test coverage is reasonably representative.




Figure 8: Fraction of Spoofer server vantage points (VPs) receiving spoofed packets of different types for the year ending August 2019, when at least one VP receives a spoofed packet. When we detect that the network hosting the client is not filtering IPv4 packets with private addresses, only a subset of the VPs receive them. Nearly all VPs receive spoofed packets of other types when the network hosting the client does not filter.

5.2 Representativeness of Spoofer data

Another inferential challenge of crowd-sourced measurement is assessing how representative the data is of the networks it samples, and the larger Internet [21]. Since the measurements we receive come from volunteers, both on-demand and triggered by clients detecting new network attachments, we do not control the sampling across either networks or time. While we can characterize the coverage of our results by, e.g., networks, ASes, or geolocation, this characterization does not capture the degree to which our measurements correctly capture the behavior of those networks.

To gain confidence in the data’s representativeness of the networks it covers, we assess the data’s predictive ability – i.e., whether we can use the data to predict the ability of a client to spoof using an IP address we have no prior tests from. We use the standard train-then-test supervised learning approach, but always split the

training and test sets such that the training samples are chronologically before the test samples. We analyze the data on month boundaries, such that the model is trained on all data prior to a given month and tested on all data for the month and after. In this way, we simulate predictions over future tests given the available data we have to that point.

We built a simple per-prefix model of spoofability that operates in ascending chronological order over training samples. Our algorithm determines the global BGP prefix from which the client ran the test using an August 2019 RIB snapshot from route-views2 [1]. We then label (or update) the prefix in a radix-trie with the client’s spoofing status. To account for variations across policies of individual networks within larger BGP prefixes, we also label the spoofing status of the more specific /24 prefix of the client. The model also computes the network-wide prior probability of spoofing using the fraction of clients that can spoof for each source address.

After building the radix-trie using the training data, the model queries the trie in a manner analogous to a routing table lookup over the test data. For each new client IP address in the test set, the model returns a prediction corresponding to the spoofing capability of its longest-prefix match on the radix trie. In the case that there is no matching longest-prefix, the model flips a biased coin that captures the overall prior probability of being able to spoof. Note that we ignore clients in the test set that also occur in the training set, so as not to artificially inflate the results positively via simple memorization of previous results from that same client IP address. Thus, we only test over clients that are “new.”

We term the ability to send spoofed packets a *positive* label. Our model’s accuracy is $> 90\%$ across IP versions and source addresses. However, accuracy is a poor measure of the model due to the fact that the prior probability of spoofability is low, so we focus on recall and precision. Figure 7 plots per-prefix spoofability recall and precision of our model, as a function of the month that splits the training from the test data. Recall is a measure of positive coverage: the fraction of the clients able to spoof for which the model correctly predicted this fact. Precision (positive predictive value) captures fidelity: the fraction of the clients the model predicted can spoof which were actually observed to spoof.

This data-driven model is basic and could incorporate many additional features, or use more advanced statistical learning techniques. Thus, our recall and precision results represent a lower-bound of what is likely possible to model. Rather, our goal is to demonstrate: i) that our automated data gathering efforts have increased the quality of the model over time; and ii) the data we have effectively captures the actual state of Internet spoofability, by virtue of its ability to make accurate predictions for unknown clients based on the prior population of client data available to the project.

Figure 7(a) shows that as the system has accumulated more tests, the model’s predictive recall has increased significantly, from a low of around 20% in May 2016, to more than 80% for IPv6 in August 2019, but only 50-70% for IPv4. Similarly, figure 7(b) shows that our model’s precision has increased significantly in the last year, and IPv6 predictions outperform IPv4 predictions. Note that 76% and 87% of the predictions come from the model’s radix trie for August 2019, for IPv4 and IPv6 respectively. Thus, the model’s performance does not come from making random guesses based on the overall prior spoofing probability, but from topological characteristics in the observed data.

The model has poor precision for IPv4 private addresses, because filtering for IPv4 private addresses is more pervasive in the core, which occludes measurement of filtering policies at the edge. Figure 8 shows that for all other classes of spoofed source address packets, nearly all Spoofer servers distributed around the world receive the class of spoofed packet. This is typically never the case with IPv4 private-address packets, which are filtered elsewhere in the network, rather than at the source.

6 NATS ARE NOT A SUBSTITUTE FOR SAV

Our testing client is primarily installed on hosts whose IPv4 connectivity is provided by a NAT [18] system. In the year ending August 2019, we received reports from 2,783 ASes using IPv4; 2,418 (86.8%) of these ASes had tests involving a NAT. A NAT system rewrites the source address of packets so that internal hosts behind the NAT with private addresses can communicate with other systems across the Internet. A well-implemented NAT should not forward packets with spoofed source addresses, as the spoofed address is unlikely to fall within the private address range the NAT serves. However, figure 4a shows that 6.4% of IPv4 /24 prefixes tested from clients behind a NAT did not filter packets with spoofed source addresses in the year ending August 2019. Further, NAT use is rare in IPv6 because unique addresses are plentiful, so SAV is explicitly required in IPv6 to filter packets with spoofed source addresses; figure 4a also shows that 12.3% of tested IPv6 /40 prefixes did not filter packets with spoofed source addresses over the same time period.

6.1 IPv4 NATs are broken

In practice, there are two failure modes where a NAT will forward packets with spoofed source addresses. We illustrate these failure modes in figure 9. In the first failure mode, the NAT simply forwards the packets with the spoofed source address (the victim) intact to the amplifier. We hypothesize this occurs when the NAT does not rewrite the source address because the address is not in the local network served by the NAT. In the second failure mode, the NAT rewrites the source address to the NAT’s publicly routable address,

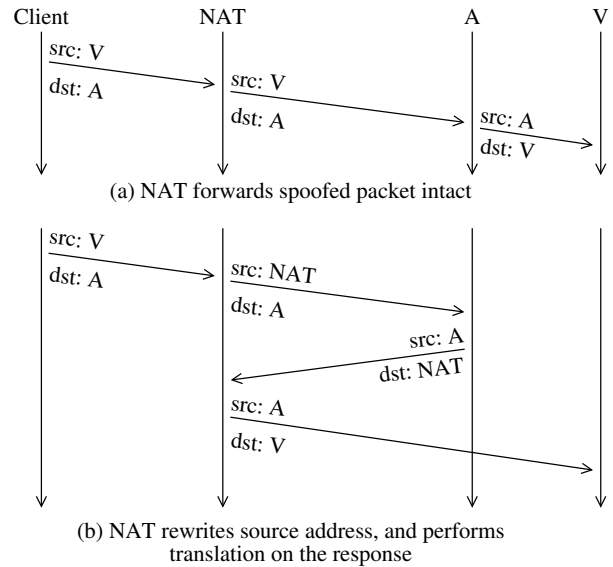


Figure 9: The two ways a NAT can forward packets with spoofed addresses. First, a NAT may forward packets with a (V)ictim’s address towards an (A)mplifier. Second, a NAT may rewrite the source address, but translate the response so the response reaches V.

and forwards the packet to the amplifier. When the server replies, the NAT system does the inverse translation of the source address, expecting to deliver the packet to an internal system. However, because the mapping is between two routable addresses external to the NAT, the packet is routed by the NAT towards the victim.

This second failure mode is important for two reasons. First, the victim still receives the attack packet, though the attacking node does not gain the benefit of amplification because it still bears the cost of sending the large packet. Second, previous studies that classified the intended victim of the attack using the source address of the packet arriving at the amplifier could have misinferred, and thus miscounted, the true victims of the attacks [12, 24]. Specifically, the source address of the packet arriving at the amplifier in figure 9b is the NAT’s external IP address, not the intended victim, who does eventually receive the amplified packet via the NAT router.

In figure 9, the amplifier and victims are addresses on separate systems, but in the Spoofer system they are assigned to the same server (figure 1) so that we can detect both failure modes. In our data collected with the probing daemon for the 11 months between September 2018 (when we began testing for the second NAT failure mode) and August 2019, we received tests from 27.8K distinct IP addresses where we detected the client was testing from behind a NAT; in comparison, we received tests from 4.6K distinct IP addresses where the client was not behind a NAT.

51.0% of NATs blocked the spoofed packets, while the remainder forwarded the packets. 46.0% of the NATs forwarded the packet after rewriting the spoofed source IP address; 3.2% of the NATs translated the destination address of our response packet back to the original spoofed address and were able to forward the response back to the Spoofer server – even though the source address (A, the

amplifier) would have caused the client’s network to discard our packet if the network had deployed SAV. 3.0% of the NATs (3.6K) forwarded the packet without rewriting the source IP address at all. In total, the Spoofer system received packets with spoofed source IP addresses from 6.2% of 27.8K IP addresses using NAT for these 11 months. In comparison, the Spoofer system received packets with spoofed source IP addresses from 13.8% of 4.6K IP addresses where the client was not behind a NAT over these same 11 months.

6.2 IPv6 looms large

While IPv6 continues to gain importance, the community’s understanding of deployed IPv6 security and hygiene practice has not kept pace [13]. IPv6 SAV is particularly important as attackers shift to leveraging new attack vectors and exploiting IoT devices, which are frequently IPv6-connected [37]. Further, IPv6 has important differences from IPv4 relating to SAV. First, whereas NATs are common in IPv4, the large address space of IPv6 means that NATs are relatively rare. By extension, the protection that should be afforded by NATs in IPv4 is missing in IPv6. Second, the immense size of the address space in IPv6 implies that attackers can utilize a much larger range of source addresses, potentially inducing state exhaustion in forwarding infrastructure.

In this subsection, we examine IPv6 SAV in detail. We first analyze filtering granularity, which is an important metric of how much of the vast IPv6 address space an attacker can spoof. Next, we infer the topological location of filtering in IPv6 as compared to IPv4, and discuss the implications for SAV deployment.

6.2.1 Filtering Granularity. A network that implements filtering to drop packets that do not originate from its prefix may still permit hosts to spoof the addresses of other hosts *within* that prefix. When SAV is in place, we term the prefix-length specificity of the policy the “filtering granularity.” Whereas we might expect more fine-grained filtering in IPv4, the large size of IPv6 assignments (even residential customers are typically allocated at least a /64 prefix that can contain 2^{64} unique hosts) suggests that within-network spoofing maybe easier. To infer filtering granularity, the client sends a series of spoofed-source packets that flip bits of the client’s true address at byte boundaries. This allows us to test the nearest adjacent address outside of a particular prefix mask.

A single Spoofer client may use many different IPv6 addresses in the same /64 prefix over time, due to the common operating system practice of using privacy extensions to assign an ephemeral lower 64bits to an IPv6 address [41]. While we may receive multiple test results from a single client with different ephemeral addresses, each result is representative of the policy of the same /64. We therefore aggregate clients into /64s in order to not bias our results.

Figure 10 shows the cumulative fraction of 15,202 unique IPv6 client /64s that could not spoof arbitrary source addresses, but could spoof within some limited range of addresses (as imposed by the filtering granularity). 70% of IPv6 clients were able to spoof sources within their /64, but not outside it. However, 21% of clients could spoof outside their /64, most commonly for /56 and /48. This matches common operator policy; some providers assign customers allocations of these common sizes [42]. In this sense, filtering in IPv6 is tighter and constrains an attacker to use addresses within




Figure 10: Inferred granularity of IPv6 filtering, restricted to clients that could not spoof arbitrarily, collected between March 2017 and August 2019. While 70% of clients could only spoof sources within their own /64, 21% could spoof outside it; ~9% were unable to spoof any address.

their own /64 prefix, which, since it is tied to the customer, does not afford anonymity or the ability to perform random source spoofing.

Interestingly, $\approx 9\%$ of clients cannot spoof at all, and are tied to their /128. All tests from the most frequent prefixes from which these clients tested were similarly constrained to use only their assigned /128 address. Note that whereas we frequently observe IPv4 clients being tied to a single address (their /32), primarily because of facilities built into DOCSIS [8], individual IPv6 client networks are frequently allocated /64 prefixes and can use any source address within that prefix.

Via personal contact, we learned that one instance of this strict IPv6 filtering is due to an OpenVPN setup. In other cases, we hypothesize that equipment strictly enforces the binding between assigned address (either via neighbor discovery or DHCP6) and source address. Indeed, we confirmed the use of “IPv6 source and prefix guard” – a mechanism to enforce exactly these bindings and prevent spoofed IPv6 traffic [11] – with one network operator.

6.2.2 Filtering Location. We examined 16,998 unique client tests and 5,687 distinct /64 IPv6 client prefixes to understand the location of SAV filtering in IPv6. These 16,998 tests come from implementing traceroute [6] on the Spoofer platform and deploying it during the period October 2017 to August 2019. Traceroute works in a fashion similar to traceroute, but reveals the location of SAV filtering.

Figure 11 shows the SAV filtering hop – relative to the source – for IPv4 with and without NAT, and for IPv6 client prefixes inferred by traceroute. While clients in 35% of IPv6 prefixes hit a filter at the first hop, most ($\approx 52\%$) were filtered at the second hop. We hypothesize that this is due to the common residential deployment scenario where the CPE is the first hop and the second hop is the provider’s router, where SAV is deployed.

While more than 50% of filtering occurs at the second hop in IPv6, 70% of IPv4 SAV is deployed at the first hop. To better understand how IPv6 and IPv4 SAV differ, we examined the difference in filter location. We found nearly 80% of the clients with IPv6 have spoofed packets filtered one hop further into the network as compared to IPv4 – likely the result of the fact that residential CPE acts as a router for IPv6, rather than a NAT. Thus, the provider’s router, which is two hops from the client, performed the filtering.




Figure 11: Inferred location of SAV filtering for clients that performed both IPv4 and IPv6 traceroute tests from October 2017 to August 2019.

6.2.3 Inter-Protocol Dependence. Finally, we explore the extent to which IPv6 SAV depends on having IPv4 SAV configured properly, and vice versa. To compute the conditional probabilities, we first find the class priors and joint probability based on frequency. We then simply divide the joint probability by the prior probability per basic probability axioms. For example the probability that IPv6 is blocked, but IPv4 is unblocked is:

$$P(v4 = received | v6 = blocked) = \frac{P(v4 = received \wedge v6 = blocked)}{P(v6 = blocked)}$$

We excluded tests where a NAT rewrote addresses, and only considered tests where a client tested both IPv4 and IPv6. We aggregated the client’s IPv4 and IPv6 addresses to their /24 and /40 prefixes, and only count each pair of IPv4, IPv6 prefix once using the most recent results for the pair. Because not only addresses, but also prefixes, are frequently dynamic in IPv6, we excluded tests where the IPv4 to IPv6 prefix mapping was not one-to-one.

Figure 12 shows the matrix of conditional probabilities for IPv4 and IPv6 filtering for the year ending August 2019. The probability that IPv6 was not filtered given that IPv4 SAV was in place is small, only 5%. However, the converse is not true; the probability that IPv4 was filtered given that IPv6 was not filtered is still 46%. And, if IPv6 was filtered, there is a high (90%) chance that IPv4 is also filtered. Interestingly, we saw a fair amount of inconsistency in IPv6 filtering when IPv4 is unfiltered; frequently IPv6 was filtered when IPv4 was not. These results suggest that operators conscientious enough to deploy SAV for IPv4 are savvy enough to also do so for IPv6, and some operators are protecting IPv6 before ensuring IPv4 is filtered. In contrast, prior work found that many IPv6 networks were relatively open relative to their IPv4 counterparts [13].

7 ANALYSIS OF REMEDIATION EFFORTS

If we receive packets with spoofed source addresses from a client at a given attachment point, but do not in subsequent tests from the same attachment point, then we infer an operator has performed remediation. Ideally, one could undertake A/B testing to measure the effect of various interventions on remediation. However, we cannot do follow-up tests ourselves; we must rely on our crowd-sourced volunteers to re-test the same network.

We have had three distinct remediation phases in the project. In May 2016, we began contacting networks where we had at least one

Probability of	v4 Filt	v4 Unfilt	v6 Filt	v6 Unfilt
v6 Unfilt	0.05	0.34	0	1
v6 Filt	0.95	0.66	1	0
v4 Unfilt	0	1	0.1	0.54
v4 Filt	1	0	0.9	0.46
	given			

Figure 12: Conditional probability between IPv4 and IPv6 SAV outcomes for clients with both address types for the year ending August 2019.

Apr 2018	NANOG (US, Canada), NLNOG (Netherlands)
May 2018	NZNOG (New Zealand), AusNOG (Australia), UKNOF (United Kingdom)
Jun 2018	GTER (Brazil)
Jul 2018	DENOG (Germany)
Sep 2018	SGOPS (Singapore)
Oct 2018	SANOG (South Asia), PacNOG (Pacific Islands), NOG.cl (Chile)
Nov 2018	JANOG (Japan), Gore (Spain)
Dec 2018	FRnOG (France)
Jan 2019	LUNOG (Luxembourg)
Feb 2019	MENOG (Middle East), ITNOG (Italy), Bolivia-NOG, ArNOG (Argentina)
Mar 2019	IDNOG (Indonesia)
Jul 2019	INNOG (India, from SANOG)

Table 3: The 20 Network Operator Groups covering 62 countries to whom we sent automatic geographically-scoped emails every month, as of August 2019.

test showing that the network did not filter packets with spoofed source addresses. We contacted the network abuse contact address registered in WHOIS, falling back to the technical contact if there was no abuse contact, or the technical contact registered in PeeringDB. We provide an example private notification email in appendix A. Then in April 2018, we began publishing geographically scoped reports of remediated and still-spoofing-capable networks to regional network operator group (NOG) email lists on a monthly basis, while continuing to privately notify networks. We provide an example public notification email in appendix B. Table 3 lists the full set we cover as of August 2019; we sent our NOG notifications using an appropriate translation for the country. Finally, we ceased sending private notifications in December 2018.

In total, we inferred 587 instances where a network hosting a Spoofer client transitioned from forwarding spoofed packets to our system, to not doing so. Figure 13 shows that 24.0% of the remediations blocking outbound spoofing occurred within a day of the first test and 35.4% within a week, indicating that an operator used our system to check for and deploy SAV. 48.2% of the remediation events we inferred took at least 1 month from the time we received the spoofed packet to when we inferred an operator had deployed SAV. Prior work observed remediation to vulnerabilities such as Heartbleed [16] occurring within a shorter period of time,




Figure 13: Delay between receiving a packet with a spoofed source address and remediation. Of the 587 remediations blocking outbound spoofing between May 2016 and August 2019, 35.4% occurred within a week, implying operators used the Spoofer client to check for and deploy SAV.

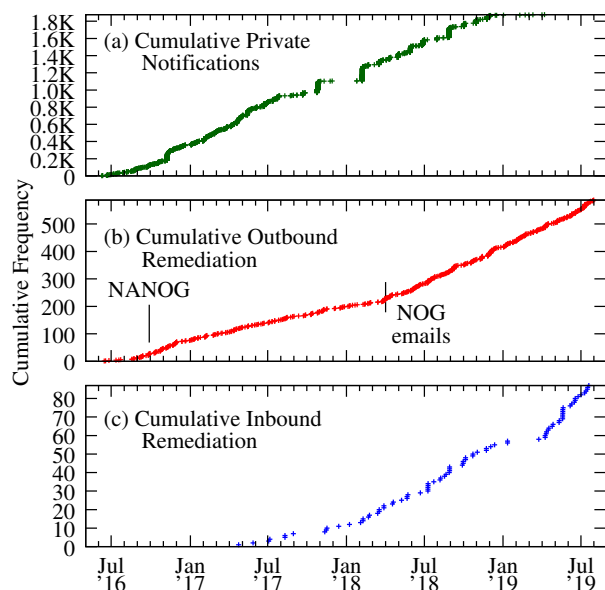


Figure 14: Cumulative private notifications and remediation inferences between May 2016 and August 2019. The number of outbound remediations we infer per month doubled when we began sending monthly NOG emails in April 2018.

consistent with the self-defense, i.e., incentive-aligned, motive for such remediation. It is difficult to separate the effect of a private notification from other forces we do not observe in the 50% of cases that took more than a month to remedy; in line with prior work, we rarely received any response to our notifications.

Figure 14a shows the cumulative private notifications we sent over time; we sent bursts of private notifications at different times (November 2016, October 2017, January 2018, and September 2018). Figure 14b shows the cumulative deployment of SAV for outbound spoofed packets over time; there were no corresponding bursts of remediation observed in figure 14b, leading us to believe the private notifications had limited impact. Further, the number of remediations we infer per month doubles from 10.6 remediations




Figure 15: Cumulative distribution of outbound remediation events per country between May 2016 and August 2019. $\approx 90\%$ of the remediation events in Brazil occurred after we began publicly sending monthly emails to GTER.

per month to 21.5 beginning April 2018 when we started publishing reports to NOG email lists, implying that public notification is more effective in promoting deployment of SAV. Figure 14c shows the cumulative deployment of SAV on inbound spoofed packets over time. We do not publicly reveal the SAV policy of individual networks on inbound packets for ethics reasons, as these represent a vulnerability to the network itself. Private notifications apparently had an impact on deploying SAV on inbound packets: when we stopped them in December 2018, the number of remediations reduced to one per month until April 2019. However, the burst of inbound remediation beginning April 2019 is not connected to any activity by the project.

Of the 587 remediation events we inferred between May 2016 and August 2019, 25.2% occurred in the U.S., and 23.5% occurred in Brazil. Figure 15 shows that nearly 90% of the remediation events in Brazil occurred after we began sending monthly emails to GTER. We calculate the remediation rate by dividing the number of ASes for which we inferred a remediation event by the total number of ASes that sent a spoofed packet during the same interval. For the year prior to commencing the GTER emails to Brazilian network operators, 14 of 67 ASes (21%) remediated; in the year after, 52 of 168 ASes (31%) remediated. This improvement is supported by NIC.br’s “Program for a Safer Internet” [45], which offers training courses and lectures to support network operators to deploy security best practices in Brazil. The rate of remediation in the U.S. is lower; prior to sending the NANOg emails to U.S. network operators, 21 of 132 (16%) of ASes remediated; in the year after, 35 of 147 (24%) of ASes remediated. While the rate of remediation is lower in the U.S. than Brazil, the relative improvement in both is equivalent $\approx 50\%$.

Figure 16 shows two metrics that reflect the effectiveness of our system at reducing the spoofed address attack surface in the Internet. The purple line is the cumulative distribution of the widest IPv4 prefix a client can spoof after remediation. Half of the remediations resulted in the client being unable to use any address apart from their single assigned address; the remainder are covered by filters that defined a range of valid addresses at the attachment point. This class of remediation represents a total of 0.13% of the total routed IPv4 address space. Because this statistic is dominated by clients only able to use their assigned address, and it is unlikely that an operator deployed SAV on a single network port, we also plot the distribution of the longest matching prefix in BGP corresponding




Figure 16: Distribution of remediation prefix length for IPv4 outbound remediations between May 2016 and August 2019. We plot the widest prefix from which the client can spoof after remediation (cumulatively representing 0.13% of the total routed address space), and the size of the longest matching prefix corresponding to the client (cumulatively, 3.01%).




Figure 17: For 4,480 unremediated IPv4/24 between May 2016 and August 2019, the distribution of unremediated network size (prefix length). We plot the size of the longest matching prefix corresponding to the client (cumulatively, 8.14% of the routed address space).

to the client’s IPv4 address, which sums to 3.01% of the total routed IPv4 address space.

While we inferred remediation in 352 IPv4/24, we found 4,480 IPv4/24 with no evidence of remediation – 12.7 times as many IPv4/24. Figure 17 shows the corresponding network sizes for the unremediated networks. Relative to figure 16, the sizes of the longest matching prefixes are similar, suggesting that the size of a prefix is not a limiting factor in the operators’ ability to deploy SAV. This space sums to 8.14% of the total routed IPv4 address space. Finally, figure 18 shows that of the /24 prefixes where we had multiple tests showing SAV had not been deployed, 21.0% had been unremediated for at least six months.

8 MOVING THE NEEDLE

Persistent lack of source address validation represents one of many failures of market forces to incentivize best security practices in the Internet ecosystem. Although the Internet engineering standards community has developed technical solutions to the spoofing vulnerability inherent in the IP architecture, gaps in SAV compliance still allow spoofed DoS attacks to grab headlines on a regular basis.




Figure 18: For 2,030 unremediated IPv4/24 with multiple positive spoofing tests between May 2016 and August 2019, the distribution of unremediated duration. 21.0% have been unremediated for at least six months.

It is clear that market forces alone will not remedy the harm that networks without SAV pose to the Internet, and to commerce that relies on it. Many efforts over the years have tried and failed to overcome the fundamental underlying problem – misaligned incentives – which hinder deployment of these technical solutions in a competitive and largely unregulated industry.

An economist would call failure to deploy SAV a *negative externality*: networks that allow spoofing save on their own operational costs, while imposing costs on others (in the form of attacks). An economic perspective argues that the only long-term remedy is to internalize this externality on the ISPs. “Naming and shaming” is a weak form of internalization. Stronger forms include liability for damages, and various types of regulation. We consider several potential future scenarios.

8.1 Impact of exogenous interventions

Section 7 concluded that our project’s approach of “naming and shaming” those who do not implement SAV had some positive impact but appears to be insufficient, based on subsequent measurements (or lack thereof) of the same networks from the Spoofer platform. But a valuable benefit of the platform is its enabling objective evaluation of the effectiveness of attempted interventions targeting remediation. We offer two examples. As of August 2019, the Internet Society had 205 distinct organizations (some with multiple ASes) participating in MANRS (§2), 159 (77.6%) asserting their commitment to SAV on the MANRS website. As part of the onboarding process, MANRS requests that the ISP send a URL showing the outcome of running Spoofer from a network without a NAT in place. For the year ending August 2019 we had IPv4 tests from 99 MANRS ASes with no NAT – likely the MANRS ISP member testing their own network, with only 11 (11.1%) able to spoof. We also had IPv4 tests from 108 MANRS ASes where a NAT was involved (more likely a representative test from a visitor to a MANRS network) and the fraction of these ASes that could spoof (25.0%) was approximately the same as the general population (22.0%, figure 4). In short, our data shows no evidence that those who assert a commitment to deploy SAV are any more likely to properly deploy it than others.

Our second example is a study of the effect of the U.S. National Science Foundation (NSF) including BCP 38 language in Campus Cyberinfrastructure (CC*) solicitations [44]. From 2014-2016, these

yearly solicitations contained language that encouraged responding institutions to address BCP 38 in their proposals: “*the Campus CI plan should address efforts to prevent IP spoofing by potential adoption of BCP 38*” and encouraged proposers to use the Spoofer system to test the current state of their network. To estimate the impact of this solicitation language, we examined tests belonging to ASes of institutions that were awarded funding from this program, and compared them to a control population of institutions who were awarded funding in political science and not in CC*. We looked for tests from an AS: i) between May 2016 and August 2019; ii) in the three-month window between the solicitation posting and the response due date (“window”); iii) anytime prior to the solicitation posting (“before”). We found that few of the institutions who received CC* funding ran the test as requested during the window of solicitation response. There were 10 awards in 2014, and 12 awards in 2016. For both the 2014 and 2016 solicitations, only two awardees in each year ran the test in the window, though 7 and 8, respectively, had run the test at some time before the solicitation. Of the 22 awardees, 3 showed evidence that they had not deployed SAV as of August 2019. In the control population, 2 of 10 awardees ran the test during the CC* solicitation window, 5 before, and none showed evidence of spoofing. We conclude that tying SAV deployment to NSF funding did not have an observable impact.

8.2 Liability, insurance, and industry standards

If network operators faced costs by assuming liability associated with attacks originating from or transiting their networks, they would have clear incentives to minimize such attacks, including by deploying technologies like SAV. Even the threat of litigation or regulation could be enough to change incentives in favor of substantially increasing SAV deployment, and might motivate insurance companies to require policy holders to provide evidence of consistent SAV deployment. As the insurance industry underwrites an increasing amount of Internet security risk, it might consider demanding SAV deployment as a way of lowering overall exposure. Inbound SAV deployment is already mandated by the widely deployed Payment Card Industry Data Security Standard (PCI DSS) [9], though §4 shows inbound SAV deployment is problematic.

Unfortunately, there are at least two stubborn barriers. The first barrier to internalizing these costs via liability for attacks is the general difficulty of attributing attacks reliably, as well as the requirement to prove economic harm. If it were feasible to attribute spoofed DoS attacks to a specific party, the associated reputational harm would already present a strong incentive to deploy SAV. A second barrier is the general presumption (enshrined in U.S. law) that networks are intermediaries who are not considered responsible for activity that merely transits their systems.

8.3 Regulating transparency

Requirements for disclosure around network management practices could serve as a stronger “name and shame” regime around SAV deployment. Such rules were part of the Federal Communications Commission (FCC) Open Internet Orders [56] and recently updated transparency requirements [57]. These requirements may already require the disclosure of SAV deployment or non-deployment, as they cover security mechanisms and device attachment rules. The

problem is likely not disclosure, but a failure of enforcement and compliance. Our tool could be an excellent arbiter of compliance with this rule, demonstrating publicly whether the network allows spoofing. This data can be useful to insurers, regulators, and to consumers wishing to understand network hygiene.

8.4 Regulating government procurement

If the U.S. Government wanted to take a leading role in increasing the ability of all networks to attribute attacks, thereby improving global cybersecurity, it could require SAV of all agency networks and require Government-contracted ISPs to support SAV as well. A similar effort successfully mandated the availability of all government websites over HTTPS with modern settings under Office of Management and Budget (OMB) Memo M-15-13 [48]. The U.S. National Institutes of Science and Technology has recently included SAV in draft security guidance documents that will represent requirements for all U.S. government agencies [43, 50]. Sometimes NIST takes these guidance documents and embeds them in Federal Information Security Modernization Act (FISMA) controls, e.g., for Domain Name System Security (DNSSEC) [59] or in other policy initiatives [60]. All such requirements are only partially effective, but they often serve as important catalysts to broader adoption.

There are several further approaches the U.S. government has still not tried: including SAV as a requirement in government-procured networking services; the Department of Homeland Security Cybersecurity and Infrastructure Security Agenda (CISA) issuing a Binding Operational Directive (BOD); or the OMB issuing a specific policy. We heard one anecdote about SAV being a requirement for Federal Risk and Authorization Management Program (FEDRAMP) technology acquisition guidelines for U.S. federal agencies, where SAV was a requirement right up until the end of the process. When the government asked for input from industry, cloud providers wanted the requirement removed because it was “too hard to implement.” This is a disturbing anecdote, since many cloud providers also sell DDoS mitigation services, so there is at least the appearance of conflict of interest in this dynamic.

This episode is reminiscent of the U.S. Anti-Bot Code (ABC) of Conduct for ISPs issued in 2012 [55]. The FCC’s Communications Security, Reliability, and Interoperability Council (CSRIC) convened a multistakeholder group to create a set of voluntary guidelines on botnet prevention and mitigation. When it was completed, the FCC asked ISPs to publicly acknowledge whether they would comply with the guidelines; the ISPs refused. This made it impossible to assess the effectiveness of the guidelines.

Two related developments challenge the prospect of increasing the strength of ISP guidelines. First, some assert that the tremendous consolidation in the Internet markets over the last twenty years has dampened the urgency of solving the SAV problem, since many companies outsource their content distribution to other platforms, e.g., one of the giant content distribution cloud platforms, many of whom have resources in place to mitigate the impact of DoS attacks by absorbing, dispersing, or blackholing attack traffic in real time [25]. Indeed, many of these cloud platforms leverage their infrastructure to sell DDoS mitigation services, so DDoS attacks represent a revenue opportunity for them. A counterpoint is that attacks are growing in volume so much that only the most

heavily capitalized providers can handle them. In October 2016, Akamai had to abandon its pro bono DDoS mitigation support for cybersecurity journalist Brian Krebs because it could not longer afford to subsidize this service. Google’s Project Shield took over Krebs’ web site instead [7]. The tremendous consolidation in interconnection may also make it easier for well-resourced networks to trace back the source of spoofed traffic as there are fewer hops to reverse engineer [10].

Second, many people tend to look at security as the responsibility of hardware and software manufacturers. In the case of the Mirai botnet [2], the U.S. Federal Trade Commission (FTC) sued the device manufacturer (D-Link) for failing to adequately secure the company’s home networking hardware [58]. We also note that a judge subsequently dismissed the lawsuit for failing to show sufficient harm by D-Link devices on consumers [54]. This does inspire the question: if a victim of a spoofed DoS attack could establish clear economic harm, and attribute it to a class of devices that did not configure SAV by default, could the equipment vendor be considered responsible for the harm?

8.5 Sticky defaults: vendor SAV responsibility

Research has found that default settings have strong impact on human behavior, even for high-stakes situations where people are well informed of their choices [23]. An important open question is why, when the benefits of deploying SAV universally are clear and the costs are low and falling, SAV is not universally deployed. Other choices of default settings in networking equipment could radically shift this equilibrium – for example, if instead of providing packets to *filter out* in network ACLs, operators had to select which packets to *forward*, they would likely make different choices and would in particular be unlikely to allow spoofed-source packets. The space of interface design for networking equipment and its impact on security is very much underexplored.

Further confirming the benefit of SAV by default is our conversations with users of the platform over the last three years, where operators think they have deployed SAV, but have not verified from all parts of their network, and since SAV is not generally a default configuration on networking equipment, pockets of spoofability can appear with any network equipment upgrade. Similarly, we have noticed many temporary conference wireless networks that support technical meetings within the Internet industry, whose operator has neglected to enable SAV when building the temporary network. While the operator often deploys SAV during the meeting after private notification, the process repeats several months later.

A related issue is network transit providers who hesitate to deploy filtering, such as with unicast Reverse Path Forwarding (uRPF) [4], because of the possibility the filtered customer network could be multihomed to another provider, now or in the future. A router that has deployed uRPF will discard a packet if the interface the packet arrived on is not the best (strict-mode) or a possible reverse path (feasible-mode) interface the router would choose to route packets to that destination. If a multihomed stub AS announces non-overlapping portions of their address space to different transit providers for traffic engineering, the provider network may find it difficult to deploy uRPF. That is, the feasible return path might not be via the interface a router received a packet from. The IETF has




Figure 19: Feasibility of uRPF over time based on observed BGP announcements across 21 years. As of August 2019, 84.3% of ASes in the Internet are Stub ASes, and 45.3% of these stub ASes in the Internet had a single inferred transit provider (38.2% of all ASes) and were candidates for feasible-mode uRPF. (Transit relationships inferred from BGP data from RouteViews and RIPE RIS using [35].)

recently proposed improvements to filtering techniques to increase their operational robustness in the face of such complexity [51].

However, we note two compelling empirical facts. First, a stub AS that is not multihomed to more than one transit provider is a candidate for at least feasible uRPF, as the transit provider will receive routes for all prefixes the stub AS uses even if the customer has multiple physical connections to their provider, or the stub AS will risk not having global connectivity in the event one connection fails. This single-homed stub AS scenario is more common than it used to be, and on the rise. Figure 19 shows that beginning 2005, as the Internet grew in terms of distinct routing policies (ASes), the trend was for stub ASes to choose a single transit provider. Transit provider ASes can deploy feasible-mode uRPF on these stub ASes without impacting packet forwarding, provided their stub AS customer properly announces prefixes covering all of their address space across each BGP session with their transit provider.

Second, more complex networks also tend to be more capitalized, and our project demonstrates (and publishes) that some of the most largest and complex networks, e.g., Comcast and AT&T, have successfully implemented SAV throughout their networks. Part of the problem, and an argument for making SAV the default, is the lack of resources (both knowledge and time) required to accurately maintain SAV filtering, confirmed in a 2017 survey of 84 operators [30]. We were gratified to hear that our platform is useful to network operators who wish to verify their own SAV compliance, including after network upgrades that created pockets of spoofability that operators did not expect. If the U.S. government mandated SAV-by-default on its networking equipment vendors, it might lead to SAV becoming the default for equipment sold into enterprise networks as well. In turn, demand for predictability by network technicians would create pressure on vendors who do not do business with the U.S. Government to make SAV a default as well.

Our data indicates that there is limited deployment of uRPF on single-homed BGP customers in the Internet. In figure 4, 25.2% of IPv4 ASes are at least partially spoofable in the year ending August 2019. For the 438 ASes where feasible-mode uRPF could be deployed that are in our data in the year ending August 2019, 21.5% of IPv4 ASes are at least partially spoofable.

9 CLOSING THOUGHTS

The Internet ecosystem, with its academic roots and radically distributed ownership, has long defied traditional governance solutions. For some vulnerabilities, there will be no simple policy solutions. For such vulnerabilities, measurement plays a critical role in quantifying the current attack surface, and assessing the effectiveness of proposed interventions. Unlike many other network security hygiene properties, there is no way to audit a network from outside to confirm that it performs SAV. The most valuable contribution of our work has been the establishment of this capability – to prove to an independent third-party auditor that one has properly deployed SAV from a given network. Any regulatory, procurement, insurance, or peering requirement would require, and thus far lacked, this measurement capability. We also validated use of this platform to fill another gap: using stored measurements to evaluate the likely effects of any deployed intervention over time. More generally, this project has been a demonstration of the importance of measurement – science, infrastructure, and data management – in developing and deploying practical solutions to the Internet’s essential security weaknesses.

ACKNOWLEDGMENTS

This material is based in part on research sponsored by the Department of Homeland Security (DHS) Science and Technology Directorate, Homeland Security Advanced Research Projects Agency, Cyber Security Division via contracts D15PC00188 and 140D7018C0010. Robert Beverly was supported in part by National Science Foundation (NSF) CNS award 1855614. The published material represents the position of the author(s) and not necessarily that of DHS or the NSF. A question from Andrei Robachevsky (ISOC) prompted us to investigate the second NAT failure mode. Gilberto Zorello (NIC.br) made us aware of BCP 38 deployment activities in Brazil. Joshua A. Kroll conducted his work while at the UC Berkeley School of Information and was supported by funding from the Berkeley Center for Law and Technology.

REFERENCES

- [1] 2019. University of Oregon RouteViews. <http://www.routeviews.org>.
- [2] Manos Antonakakis, Tim April, Michael Bailey, Matt Bernhard, Elie Bursztein, Jaime Cochran, Zakir Durumeric, J. Alex Halderman, Luca Invernizzi, Michalis Kallitsis, Deepak Kumar, Chaz Lever, Zane Ma, Joshua Mason, Damian Menscher, Chad Seaman, Nick Sullivan, Kurt Thomas, and Yi Zhou. 2017. Understanding the Mirai Botnet. In *26th USENIX Security Symposium*. Vancouver, BC, 1093–1110.
- [3] Arbor Networks. 2018. Worldwide Infrastructure Security Report. <https://resources.arbornetworks.com/>.
- [4] F. Baker and P. Savola. 2004. *Ingress Filtering for Multihomed Networks*. RFC 3704.
- [5] Robert Beverly and Steven Bauer. 2005. The spoofer project: inferring the extent of source address filtering on the internet. In *Proceedings of the USENIX Steps to Reducing Unwanted Traffic on the Internet Workshop (SRUTI)*.
- [6] Robert Beverly, Arthur Berger, Young Hyun, and kc claffy. 2009. Understanding the efficacy of deployed Internet source address validation filtering. In *IMC*.
- [7] Brian Krebs. 2016. <https://www.zdnet.com/article/krebs-on-security-booted-off-akamai-network-after-ddos-attack-proves-pricey/>.
- [8] Cablelabs. 2006. Data Over Cable Service Interface Specification (DOCSIS). <http://www.cablemodem.com/>.
- [9] Alan Calder and Geraint Williams. 2015. *PCI DSS: A Pocket Guide* (4th ed.). It Governance Ltd.
- [10] Yi-Ching Chiu, Brandon Schlinker, Abhishek Balaji Radhakrishnan, Ethan Katz-Bassett, and Ramesh Govindan. 2015. Are We One Hop Away from a Better Internet?. In *IMC*.
- [11] Cisco. 2019. IPv6 Source Guard and Prefix Guard. [- \[src-guard.html\]\(#\).
 - \[12\] Jakub Czym, Michael Kallitsis, Manaf Gharaibeh, Christos Papadopoulos, Michael Bailey, and Manish Karir. 2014. Taming the 800 Pound Gorilla: The Rise and Decline of NTP DDoS Attacks. In *IMC*. 435–448.
 - \[13\] Jakub Czym, Matthew Luckie, Mark Allman, and Michael Bailey. 2016. Don’t Forget to Lock the Back Door! A Characterization of IPv6 Network Security Policy. In *NDSS*.
 - \[14\] L. Daigle. 2004. *WHOIS Protocol Specification*. RFC 3912.
 - \[15\] Z. Duan, X. Yuan, and J. Chandrashekar. 2006. Constructing Inter-Domain Packet Filters to Control IP Spoofing Based on BGP Updates. In *INFOCOM*.
 - \[16\] Zakir Durumeric, James Kasten, David Adrian, J. Alex Halderman, Michael Bailey, Frank Li, Nicholas Weaver, Johanna Amann, Jethro Beekman, Mathias Payer, and Vern Paxson. 2014. The Matter of Heartbleed. In *IMC*.
 - \[17\] W. Eddy. 2007. *TCP SYN Flooding Attacks and Common Mitigations*. RFC 4987.
 - \[18\] K. Egevang and P. Francis. 1994. *The IP Network Address Translator \(NAT\)*. RFC 1631.
 - \[19\] P. Ferguson and D. Senie. 2000. *Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing*. RFC 2827.
 - \[20\] Marcella Hastings, Joshua Fried, and Nadia Heninger. 2016. Weak keys remain widespread in network devices. In *IMC*. 49–63.
 - \[21\] Gokay Huz, Steven Bauer, kc claffy, and Robert Beverly. 2015. Experience in using Mechanical Turk for Network Measurement. In *Proceedings of the ACM SIGCOMM Workshop on Crowdsourcing and Crowdfunding of Big Internet Data*.
 - \[22\] Internet Society. 2019. Mutually Agreed Norms for Routing Security \(MANRS\). <https://www.manrs.org/>.
 - \[23\] Eric J Johnson and Daniel Goldstein. 2003. Do defaults save lives? *Science* 302, 5649 \(Nov. 2003\), 1338–1339.
 - \[24\] Mattijs Jonker, Alistair King, Johannes Krupp, Christian Rossow, Anna Sperotto, and Alberto Dainotti. 2017. Millions of Targets Under Attack: a Macroscopic Characterization of the DoS Ecosystem. In *IMC*.
 - \[25\] K. Claffy and D. Clark. 2019. Workshop of Internet Economics 2019 Final Report. <https://www.caida.org/outreach/workshops/wie/>.
 - \[26\] Sam Kottler. 2018. GitHub Engineering DDoS Incident Report. <https://githubengineering.com/ddos-incident-report/>.
 - \[27\] Brian Krebs. 2016. The Democratization of Censorship. <https://krebsonsecurity.com/2016/09/the-democratization-of-censorship/>.
 - \[28\] Frank Li, Zakir Durumeric, Jakub Czym, Mohammad Karami, Michael Bailey, Damon McCoy, Stefan Savage, and Vern Paxson. 2016. You’ve Got Vulnerability: Exploring Effective Vulnerability Notifications. In *USENIX Security Symposium*. 1033–1050.
 - \[29\] Frank Li, Grant Ho, Eric Kuan, Yuan Niu, Lucas Ballard, Kurt Thomas, Elie Bursztein, and Vern Paxson. 2016. Remediating Web Hijacking: Notification Effectiveness and Webmaster Comprehension. In *WWW*. 1009–1019.
 - \[30\] Franziska Lichtblau, Florian Streibelt, Thorben Krüger, Philipp Richter, and Anja Feldmann. 2017. Detection, Classification, and Analysis of Inter-domain Traffic with Spoofed Source IP Addresses. In *IMC*.
 - \[31\] B. Liu, J. Bi, and A. V. Vasilakos. 2014. Toward Incentivizing Anti-Spoofing Deployment. *IEEE Transactions on Information Forensics and Security* 9, 3 \(March 2014\), 436–450.
 - \[32\] Xin Liu, Ang Li, Xiaowei Yang, and David Wetherall. 2008. Passport: Secure and Adoptable Source Authentication. In *NSDI*. USENIX.
 - \[33\] Qasim Lone, Matthew Luckie, Maciej Korczyński, Hadi Asghari, Mobin Javed, and Michel van Eeten. 2018. Using Crowdsourcing Marketplaces for Network Measurements: The Case of Spoofer. In *Network Traffic Measurement and Analysis Conference \(TMA\)*. 1–8.
 - \[34\] Qasim Lone, Matthew Luckie, Maciej Korczyński, and Michel van Eeten. 2017. Using loops observed in Traceroute to infer the ability to Spoof. In *Passive and Active Network Measurement Conference \(PAM\)*. 229–241.
 - \[35\] Matthew Luckie, Bradley Huffaker, Amogh Dhamdhare, Vasilieos Giotsas, and kc claffy. 2013. AS Relationships, Customer Cones, and Validation. In *IMC*. 243–256.
 - \[36\] Sean Lyngaas. 2019. Someone is spoofing big bank IP addresses – possibly to embarrass security vendors. <https://www.cybercoop.com/spoofed-bank-ip-address-greynoise-andrew-morris-bank-of-america/>.
 - \[37\] Kieren McCarthy. 2018. It’s begun: ‘First’ IPv6 denial-of-service attack puts IT bods on notice. \[https://www.theregister.co.uk/2018/03/03/ipv6_ddos/\]\(https://www.theregister.co.uk/2018/03/03/ipv6_ddos/\).
 - \[38\] Doug Montgomery and Kotikalapudi Sriram. 2017. Evaluation and Deployment of Advanced DDoS Mitigation Techniques. <https://www.nist.gov/sites/default/files/documents/2017/09/22/ddosd-nist-2016-08-v3.pdf>.
 - \[39\] Christopher Morrow. 2006. BLS FastAccess internal tech needed. <http://www.merit.edu/mail.archives/nanog/2006-01/msg00220.html>.
 - \[40\] T. Mrugalski, M. Siodelski, B. Volz, A. Yourtchenko, M. Richardson, S. Jiang, T. Lemon, and T. Winters. 2018. *Dynamic Host Configuration Protocol for IPv6 \(DHCPv6\)*. RFC 8415.
 - \[41\] T. Narten, R. Draves, and S. Krishnan. 2007. *Privacy Extensions for Stateless Address Autoconfiguration in IPv6*. RFC 4941.
 - \[42\] T. Narten, G. Huston, and L. Roberts. 2011. *IPv6 Address Assignment to End Sites*. RFC 6177.](https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6_fhsec/configuration/xs-3s/ipv6f-xe-3s-book/ipv6-</div><div data-bbox=)

- [43] National Institutes of Standards and Technology. 1990. NIST Special Publication 800-series General Information. <https://www.nist.gov/itl/nist-special-publication-800-series-general-information>.
- [44] National Science Foundation. 2016. Campus Cyberinfrastructure (CC*), NSF 16-567. <https://www.nsf.gov/pubs/2016/nsf16567/nsf16567.htm>.
- [45] NIC.br. 2018. Program for a Safer Internet. <https://bcp.nic.br/i+seg/>.
- [46] Zhiyun Qian, Z. Morley Mao, Yinglian Xie, and Fang Yu. 2010. Investigation of Triangular Spamming: A Stealthy and Efficient Spamming Technique. In *IEEE Security and Privacy*.
- [47] Stefan Savage, David Wetherall, Anna Karlin, and Tom Anderson. 2000. Practical Network Support for IP Traceback. In *SIGCOMM*. 295–306.
- [48] Tony Scott. 2015. M-15-13: Policy to Require Secure Connections across Federal Websites and Web Services. <https://https.cio.gov/>.
- [49] Alex C. Snoeren, Craig Partridge, Luis A. Sanchez, Christine E. Jones, Fabrice Tchakountio, Stephen T. Kent, and W. Timothy Strayer. 2001. Hash-based IP Traceback. In *SIGCOMM*. 3–14.
- [50] Kotikalapudi Sriram and Doug Montgomery. 2018. Secure Interdomain Traffic Exchange - BGP Robustness and DDoS Mitigation. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP) 800-189. <https://src.nist.gov/publications/detail/sp/800-189/draft>.
- [51] K. Sriram, D. Montgomery, and J. Haas. 2019. Enhanced Feasible-Path Unicast Reverse Path Filtering. <https://tools.ietf.org/html/draft-ietf-osepc-urpf-improvements-03>.
- [52] Ben Stock, Giancarlo Pellegrino, Christian Rossow, Martin Johns, and Michael Backes. 2016. Hey, You Have a Problem: On the Feasibility of Large-Scale Web Vulnerability Notification. In *USENIX Security Symposium*. 1015–1032.
- [53] US-CERT. 2008. Multiple DNS implementations vulnerable to cache poisoning VU#800113.
- [54] U.S. District Court of Northern District of California. 2017. Order Re: Motion to Dismiss: Federal Trade Commission v. D-link Systems. <https://consumerist.com/consumermediallc.files.wordpress.com/2017/09/dlinkdismissal.pdf>.
- [55] U.S. Federal Communication Commission. 2012. Final Report: U.S. Anti-Bot Code of Conduct (ABCs) for Internet Service Providers (ISPs). <https://transition.fcc.gov/bureaus/pshs/advisory/csric3/CSRIC-III-WG7-Final-Report.pdf>.
- [56] U.S. Federal Communication Commission. 2015. In the Matter of Protecting and Promoting the Open Internet: Report and Order on Remand, Declaratory Ruling, and Order. 30 FCC Rcd 5601 (7); 80 FR 19737.
- [57] U.S. Federal Communication Commission. 2018. Restoring Internet Freedom. 33 FCC Rcd 311 (1).
- [58] U.S. Federal Trade Commission. 2017. FTC Charges D-Link Put Consumers' Privacy at Risk Due to the Inadequate Security of Its Computer Routers and Cameras. <https://www.ftc.gov/news-events/press-releases/2017/01/ftc-charges-d-link-put-consumers-privacy-risk-due-inadequate>.
- [59] U.S. Government Office of Management and Budget. 2008. Securing the Federal Government's Domain Name System Infrastructure. <https://www.whitehouse.gov/sites/whitehouse.gov/files/omb/memoranda/2008/m08-23.pdf>.
- [60] U.S. Government Office of Management and Budget. 2010. Internet Protocol Version 6 (IPv6), mandated transition. https://obamawhitehouse.archives.gov/sites/default/files/omb/assets/egov_docs/transition-to-ipv6.pdf.
- [61] Nick Weaver. 2006. Can You Spoof IP Packets? <https://slashdot.org/story/06/05/02/1729257/can-you-spoof-ip-packets>.
- [62] A. Yaar, A. Perrig, and D. Song. 2006. StackPi: New Packet Marking and Filtering Mechanisms for DDoS and IP Spoofing Defense. *IEEE JSAC* 24, 10 (Oct 2006).

A PRIVATE NOTIFICATION

The Spoofer system automatically generated emails for individual ASes using the following general format. The system tailored the details of the notification to the AS using the evidence it had available. For example, it included a secret key in the URL so that the contact could view results of inbound spoofing tests. We manually sent each email to the abuse or other technical contact using the real name and contact address of a representative of the project.

From: Matthew Luckie <mjl@caida.org>
To: <abuse contact>
Subject: source IP address spoofing from <name of network>

While reviewing recent public tests from the CAIDA spoofer client

<https://www.caida.org/projects/spoofer/>

I came across one involving <name of network>. It seems that based on the testing history for AS<num>, there is inadequate filtering of IPv6 packets with invalid source addresses, so packets with spoofed IPv6 source addresses can leave your network. These systems can participate in volumetric denial of service attacks. However, it seems that packets with spoofed source IPv4 addresses are correctly being filtered. Further, packets with spoofed source addresses claiming to be from inside your network are not filtered when they arrive from outside your network.

https://spoofer.caida.org/recent_tests.php?as_include=<num>

<https://www.ietf.org/rfc/rfc2827.txt>

B PUBLIC REGION-FOCUSED NOTIFICATION

The Spoofer system automatically generated region-focused emails and sent them using a role account to network operator group mailing lists, using the following general format. The system sorted the ASes in the improvements table by the date it inferred remediation to take place, and sorted ASes in the issues table by the date it first observed spoofed packets. We used translations for the report in French, Indonesian, Italian, Japanese, Portuguese, and Spanish.

From: CAIDA Spoofer Project <spoofer-info@caida.org>

To: <NOG email list>

Subject: CAIDA Spoofer Report for <NOG> for <month>

In response to feedback from operational security communities, CAIDA's source address validation measurement project (<https://spoofer.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 summarizes tests conducted within <countries>

Inferred improvements during <month>:

ASN	Name	Fixed-By
64496	IANA-RSVD #1	<yyyy-mm-dd>
64497	IANA-RSVD #2	<yyyy-mm-dd>

Further information for the inferred remediation is available at: <https://spoofer.caida.org/remedy.php>

Source Address Validation issues inferred during <month>:

ASN	Name	First-Spoofed	Fixed-By
64498	IANA-RSVD #3	<yyyy-mm-dd>	<yyyy-mm-dd>
64499	IANA-RSVD #4	<yyyy-mm-dd>	<yyyy-mm-dd>
64500	IANA-RSVD #5	<yyyy-mm-dd>	<yyyy-mm-dd>

Further information for these tests where we received spoofed packets is available at:

https://spoofer.caida.org/recent_tests.php?no_block=1&country_include=<countries>

Please send any feedback or suggestions to spoofer-info@caida.org