Augment Spoofer Project to Improve Remediation Efforts (ASPIRE) Final Report

Dr. Kimberly Claffy and Dr. Matthew Luckie

September 21, 2020

Table of Contents

A Executive Summary

1. Performer: University of California, San Diego (UCSD) & University of Waikato, NZ
2. Award: DHS S&T contract 140D7018C0010.
3. Period of Performance: 1 Sept 2018 - 19 Sept 2020.

Despite source IP address spoofing being a known vulnerability – arguably the greatest architectural vulnerability in the TCP/IP protocol suite as designed – for close to 30 years, and despite many efforts to shed light on the problem, spoofing remains a viable attack method for redirection, amplification, and anonymity. While some application-layer patches can mitigate these attacks, attackers continuously search for new vectors. To defeat DDoS attacks requires operators to ensure their networks filter packets with spoofed source IP addresses, a *best current practice (BCP)* known as source address validation (SAV). The overarching objective of our project was to promote using SAV BCP by networks around the world, addressing a long-standing threat to U.S. critical communciations infrastructure.

This final report describes the following accomplished deliverables, also linked from CAIDA's web site.[1] (1) updates to the open source client-server source address validation (SAV) testing system (developed under DHS S&T contract D15PC00188) to expand visiblity of networks behind Network Address Translation devices (NATs); (2) expanded notifications and reporting through our operator-focused private reporting engine and public regionally-focused notifications to operational mailing lists; (3) documente analysis of the effectiveness of different approaches to stimulating remediation activities. These tasks achieved testing and evaluation of work developed under the previous contract, and technology transition to a broader cross-section of security research, operations, commercially provided risk management ecosystem, and public policy stakeholders. The resulting technologies and data improved the U.S. government's ability to identify, monitor, and mitigate the infrastructure vulnerability that serves as the primary vector of massive DDoS attacks on the Internet.

This project exemplifies how the Internet ecosystem has long defied traditional governance solutions. For some cyberinfrastructure 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 anti-spoofing source address validation (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 the importance of measurement – science, infrastructure, and data management – in developing and deploying practical solutions to the Internet's essential security weaknesses.

[1] https://www.caida.org/funding/ddos-aspire/.

B Problem: Long-standing vulnerability that threatens U.S. critical infrastructure

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 [21, 23] (DoS), resource exhaustion [17], policy evasion [34], and cache poisoning [43] 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 [32] 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. To defeat spoofed-source DDoS attacks requires operators to ensure their networks filter packets with spoofed source IP addresses [18], a *best current practice (BCP)* known as source address validation (SAV). However, a network's deployment of SAV primarily helps other networks, and is categorically incentive-incompatible, since a mistake configuring SAV or failure to keep it current could accidentally discard valid customer packets. SAV represents a classic tragedy of the commons in the Internet.

Recognizing the challenge of regulation in this area, many public and private sector efforts have experimented with techniques to promote deployment of SAV [15, 51, 27, 26]. Also, in 2014, the Internet Society began to foster grassroots community support to launch the global MANRS initiative – Mutually Agreed Norms for Routing Security [19], which included a public commitment to deploy source address validation, among other routing security best practices. 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 [33], and in 2018 provided deployment guidance [42].

However, measuring the effectiveness of any of these approaches presents its own challenges. Unlike other best practices which trusted third parties can scan to verify compliance with, verifying a network's compliance with SAV requires a measurement vantage point *inside* (or immediately upstream of) that network, because the origin network of arbitrary spoofed packets cannot be determined [2]. During the previous three year contract (DHS S&T contract D15PC00188), we built a production-quality software client that volunteers across the Internet could download and run from networks they visit, testing those networks' ability to send various types of spoofed packets to our server, which collects, aggregates, and reports test results.

C Technical Approach

Our system architecture includes: (1) a server instance that coordinates measurements and obtains results, (2) client software with a graphical user interface for Windows, MacOS, and UNIX-like systems, and (3) a set of distributed Ark nodes that receive spoofed packets and allow us to infer

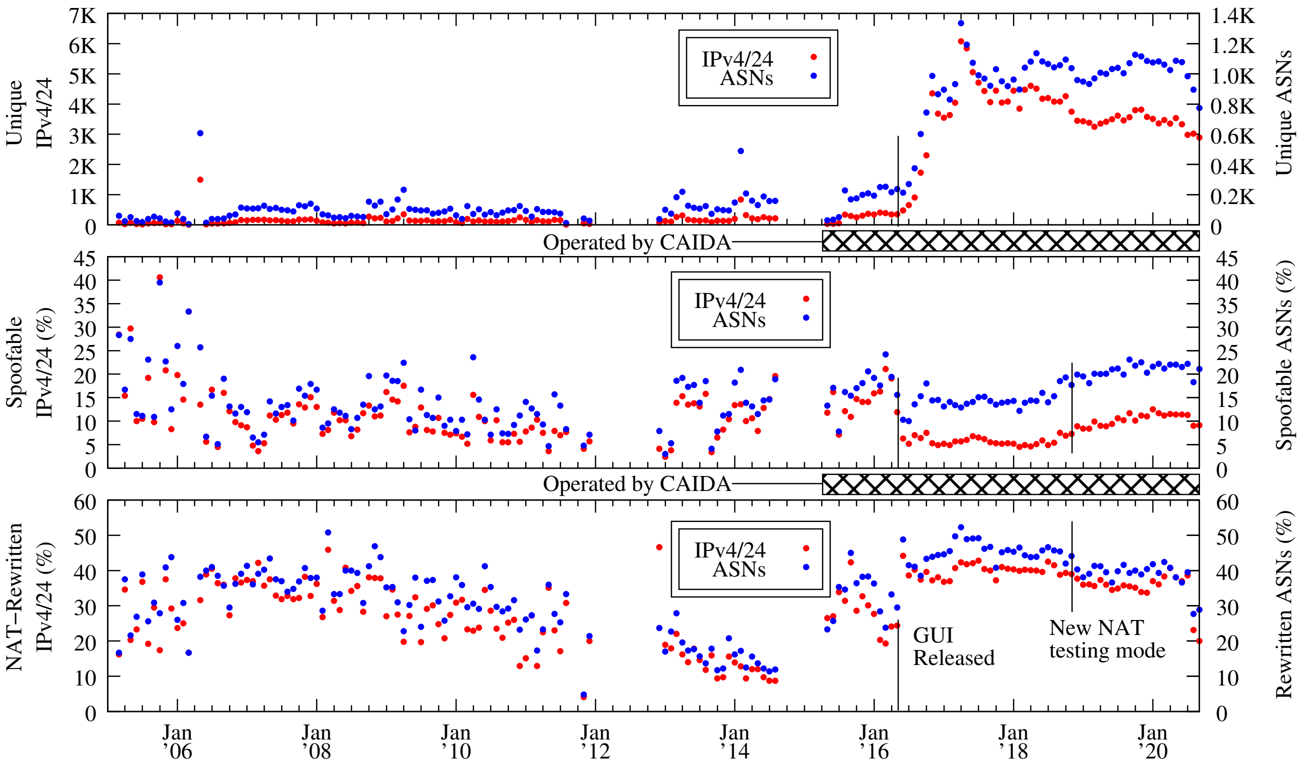

Figure 1: *Overview of Spoofer project data collection over time, aggregated per month. The gaps prior to May 2015 are due to hardware failures. After we released our new software system in May 2016 (during the previous project DHS D15PC00188), the volume of tests increased an order of magnitude from ≈300 IPv4 /24 prefixes in ≈200 ASes to ≈4K IPv4 prefixes in ≈1K ASes per month. Between November 2016 and June 2018, the range of spoofable IPv4 prefixes was 4.9% – 6.8%, and the range of spoofable ASNs was 13.1% – 15.1%. However, the range of prefixes with tests that reveal rewritten source addresses over this same period was 37.2% – 42.8%, and the range of ASNs with tests with rewritten source addresses 43.8% – 52.3%; these clients represented a gap in our visibility into SAV deployment. We developed a new measurement technique (Figure 3) to close this gap, allowing us to target remediation more precisely.*

where along a path SAV may be configured.[2] The resulting data is used to inform the operational security community regarding which networks permit spoofed packets to exit their networks.

For this new project, we delivered four complementary efforts: improving capabilities of the Spoofer software to close visibility gaps with new measurement and visualization techniques; implementation and analysis of methods to improve remediation outcomes, including expanded reporting functionality; investigation of integration of data products and capabilities into the commercially-provided risk management ecosystem; and maintaining the production-quality client/server system to ensure continued data acquisition. Figure 1 provides an overview of the Spoofer Project's data collection since the project began in 2005. (DHS funding began in August 2015.)

Figure 2 provides an overview of our findings on SAV deployment, for packets both outbound

[2] The Ark infrastructure is supported by funding from the National Science Foundation (NSF OAC-1724853) and DHS Cooperative Agreement FA8750-18-2-0049.

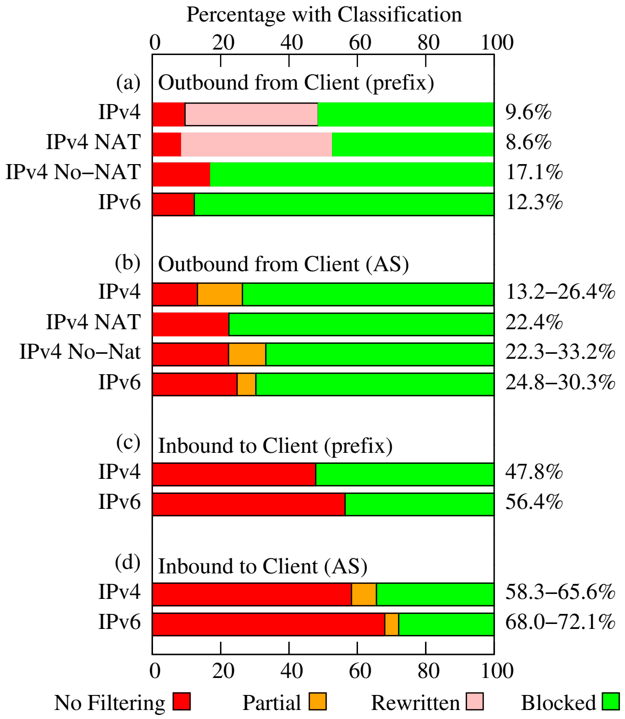

Figure 2: *Summary statistics for Spoofer Project for the year ending August 2020. Percent values to right of barplot indicate the fraction of networks with problematic SAV deployments, either no filtering or partial filtering of spoofed packets (red plus orange segments).*

from and inbound to the measured network, for the year ending 1 August 2020. 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. For example, 26.4% and 30.3% 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 2b). The comparatively small fraction of prefixes from which we received spoofed IPv4 packets (figure 2a) 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, 17.1% of IPv4 prefixes had no filtering; filtering was better deployed at prefix-level granularity in IPv6, with 12.3% of tested prefixes not filtering.

The lower panels of figure 2 (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. 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 [39], thus treating externally-sourced packets as internal represents a risk to one's own network. In our data, 65.6% and 72.1% of IPv4 and IPv6 ASes, respectively, had at least one prefix that was not filtering inbound packets.

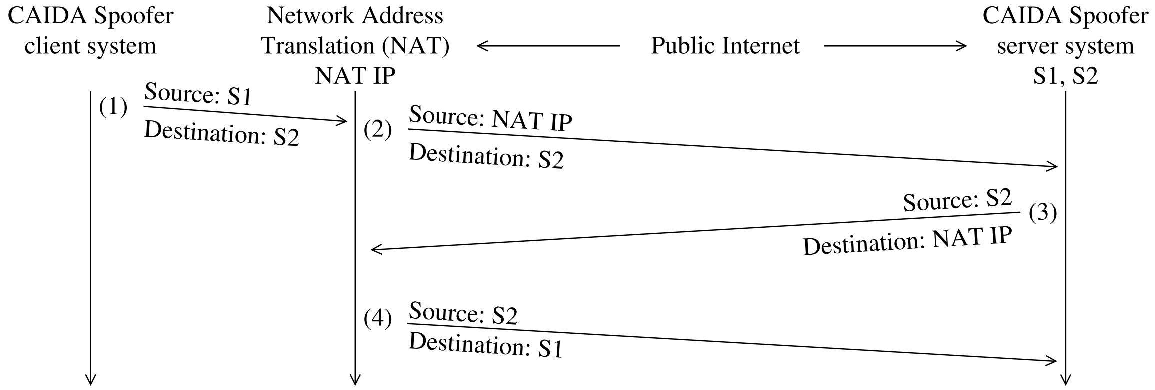

Figure 3: *We extended the Spoofer client-server system to test clients whose packets are rewritten with spoofed source addresses. Our system sends packets (1), and classifies the client's spoofed packets as rewritten because source address S1 in packet (1) is changed to the NAT router's public IP address in packet (2). However, if we send a response to the rewritten packet as in packet (3), some NAT routers will respond with packet (4) to the Spoofer server when they translate the source address of packet (3) to the NAT IP to S2 in packet (4), outbound from the NAT, even though S2 is not an internal address of the network the router is attached to. If we receive packet (4) we can infer the client's network does not perform source address validation.*

D Development Deliverables

D.1 Extend visibility capabilities to infrastructure behind NATs

We extended the capabilities of our Spoofer software to further test networks whose Network Address Translation (NAT) router obscures our view of SAV deployment. We modified our server software to respond to packets whose spoofed source address has been rewritten to infer SAV policy for these networks.

The critical element of our approach is as follows. We discovered that if we send a response to the rewritten packet, some NAT routers will translate the responding source address of the packet back to the original spoofed source address that we control, and then forward the packet to us. Figure 3 illustrates this scenario. We first instruct the client system to use S1 as its spoofed source address in a packet the client sends to S2, and both of these addresses we have assigned to the Spoofer server. The NAT router will rewrite the source address of the packet to the NAT IP, and forward the packet to the Spoofer server. We then send a reply to the NAT router, which will perform the inverse translation, i.e. swapping NAT IP back to the original source IP address S2. If we receive a responding packet that has S2 as its *source*, we can infer that the network has not deployed SAV, and use that indication to trigger remediation efforts.

Figure 1 shows a vertical line that shows when we activated this new NAT testing mode. As our new NAT testing mode allows greater visibility into networks that do not block spoofed packets, the fraction of networks we classify as rewriting the spoofed probe dropped, and the fraction of networks we classify as spoofable increased.

D.2 Support for AS-level opt-in notification system

When we receive a test showing spoofed packets are not blocked, our current approach to operator notification is to privately notify the abuse contact recorded for the AS in the WHOIS database, or a technical contact for the AS in PeeringDB [38]. However, these notifications do not necessarily reach the appropriate technical contact within an AS who can effect remediation. Representatives in the operational security community have requested that we automatically notify them should we ever receive a spoofed packet from their network. Therefore, we created a registration system in the Spoofer project's reporting engine that allows a vetted operator within an AS to receive these notifications. In the event that we receive a test showing a prefix without source address validation deployed, instead of sending emails to the WHOIS-registered abuse contact, or a technical contact recorded in PeeringDB, we send the email to the registered contact.[3]

D.3 Automated monthly reports to regional operator mailing lists

Representatives in the operational security community encouraged us to begin sending public notifications where we infer SAV is not deployed. Beginning March 2018, we have sent monthly emails to numerous public region-focused network operator group (NOG) emailing lists:

1. AONOG – Angola
2. AusNOG – Australia [3]
3. ATNOG – Austria
4. DENOG – Germany
5. ESNOG – Spain, which we translated to Spanish
6. FRnOG – France, which we translated to French
7. GTER – Brazil, which we translated to Portuguese [6]
8. INNOG – India
9. ITNOG – Italy, which we translated to Italian
10. JANOG – Japan, which we translasted to Japanese
11. LUNOG – Luxembourg
12. MENOG – Bahrain, Egypt, Iran, Iraq, Jordan, Kuwait, Lebanon, Oman, Palestinian Territory, Qatar, Saudi Arabia, Syria, Turkey, United Arab Emirates, and Yemen
13. PacNOG – American Samoa, Cook Islands, Fiji, Federated States of Micronesia, Guam, Kiribati, Marshall Islands, Northern Mariana Islands, New Caledonia, Niue, Nauru, Palau, Papua New Guinea, French Polynesia, Solomon Islands, Tokelau, Tonga, Tuvalu, Vanuatu, Wallis and Futuna Islands, Samoa
14. NANOG – the United States and Canada [7]
15. NOG Bolivia – Bolivia
16. NOG.cl – Chile
17. NLNOG – the Netherlands [4]
18. NZNOG – New Zealand [5]
19. SANOG – Afghanistan, Bangladesh, Bhutan, Sri Lanka, Maldives, Nepal, Pakistan
20. SGNOG – Singapore
21. UKNOF – the United Kingdom [8]

[3] Operators can register for this system at https://spoofer.caida.org/register.php.

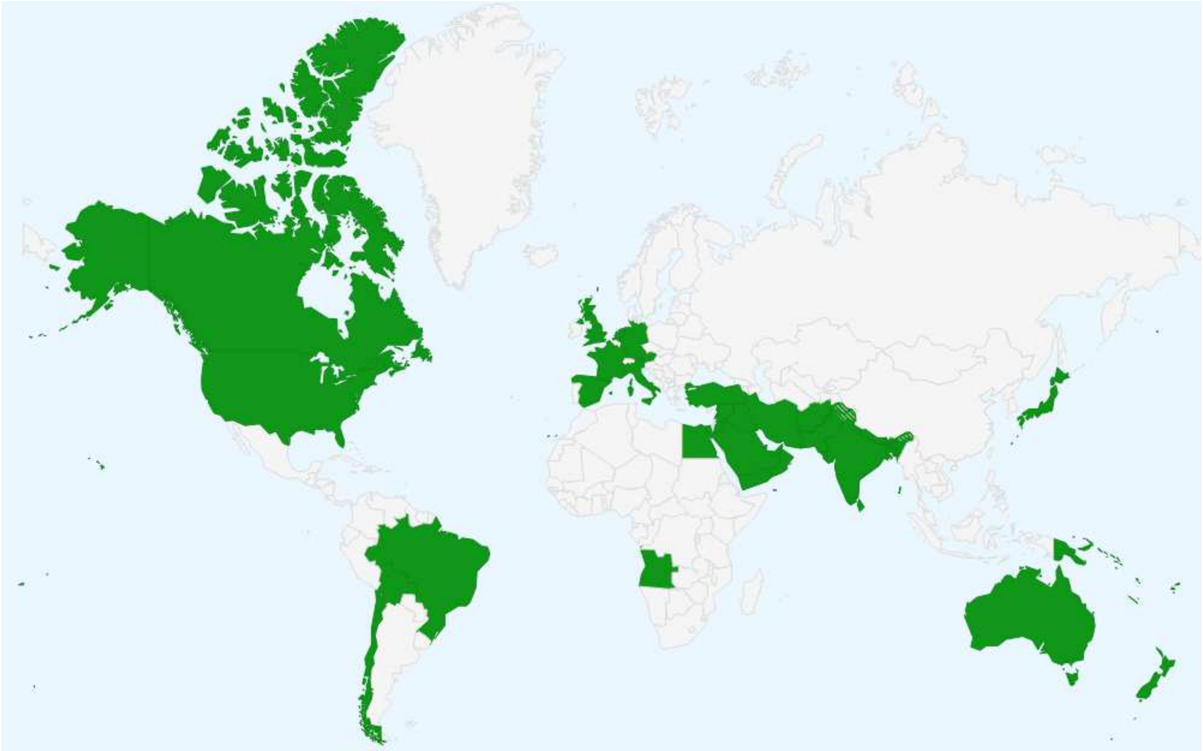

Figure 4: *This map shows the countries currently covered by our monthly emails to public region-focused network operator group (NOG) emailing lists in September 2020.*

Figure 4 colors in green countries we cover with monthly reports. Prior to sending the first monthly report, we received permission from the NOG email list administrators to send these emails. Our emails report ASes within each region where we infer remediation activity has taken place, as well as ASes originating prefixes from which we have received spoofed packets in the previous month. We received public support for this activity from the Internet Society (ISOC) Mutually Agreed Norms for Routing Security (MANRS) initiative [24]. We used our network of international collaborators to expand these public notifications to include additional countries, including translating the reports into the native language spoken in those countries where possible.

D.4 Maintain spoofer operations

We continued to support the Spoofer system on modern platforms, ensuring our spoofer software continues to work with new operating system releases for Windows, MacOS, and Linux. We upgraded the operating system on the deployed spoofer servers to ensure the underlying software is still supported by the vendor. We expanded our efforts to build Spoofer packages for home access router platforms, e.g., OpenWRT-based, and investigated options for integration of the software into other home access router platforms. We implemented an improved visual interface to the data, illustrated in Figure 5.

Elements of the improved AS-level graph in a Spoofer Session Report

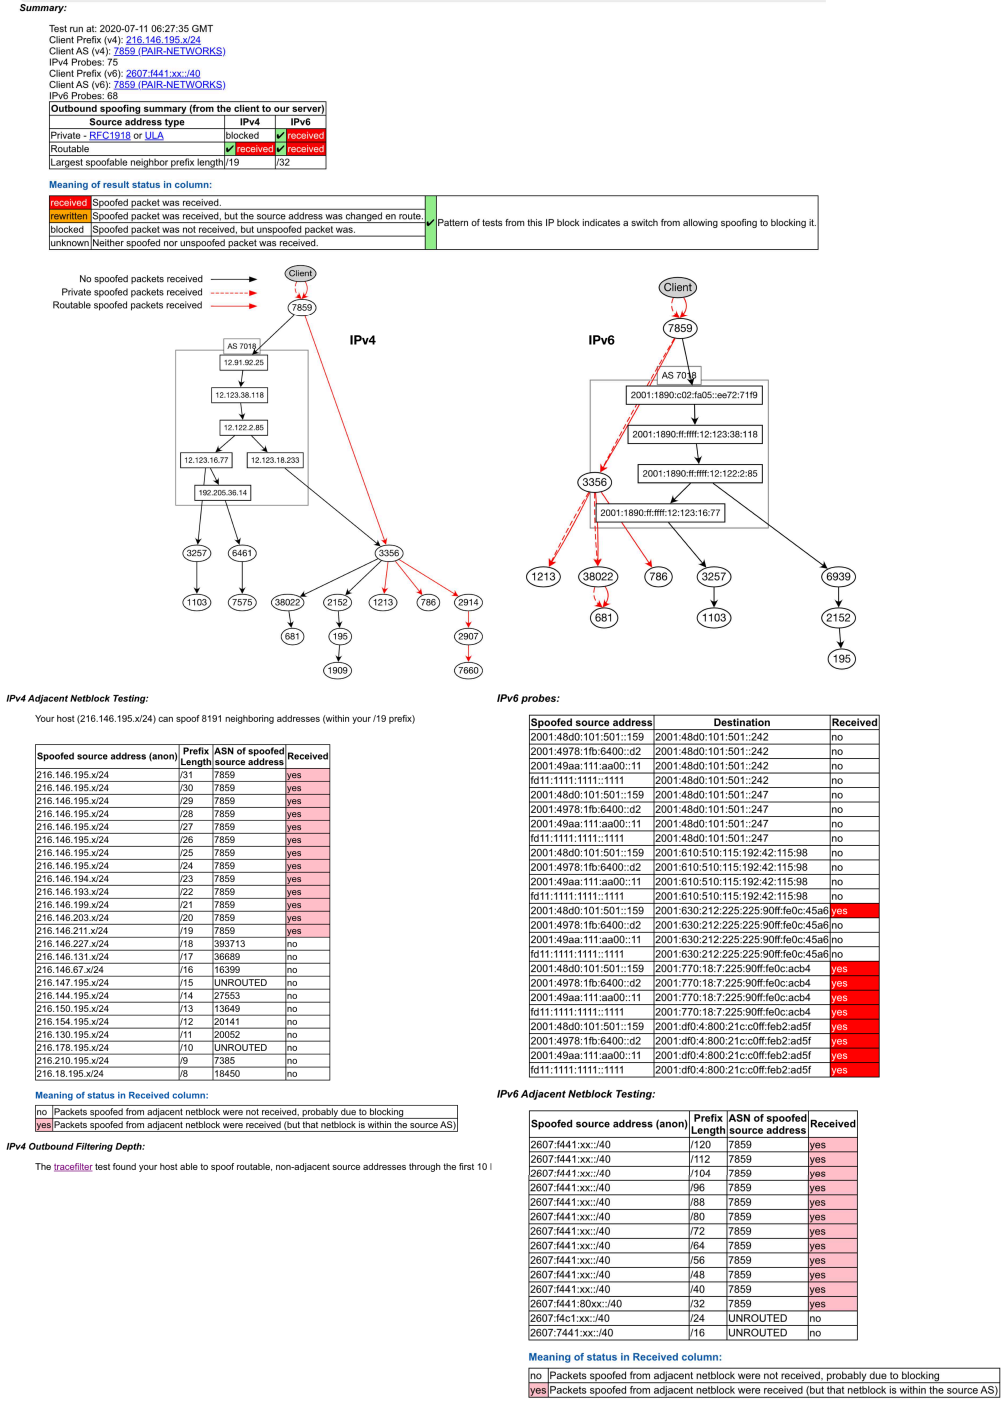

Summary:

Test run at: 2020-07-11 06:27:35 GMT
Client Prefix (v4): 216.146.195.x/24
Client AS (v4): 7859 (PAIR-NETWORKS)
IPv4 Probes: 75
Client Prefix (v6): 2607:f441:xx::/40
Client AS (v6): 7859 (PAIR-NETWORKS)
IPv6 Probes: 68

Outbound spoofing summary (from the client to our server)		
Source address type	**IPv4**	**IPv6**
Private - RFC1918 or ULA	blocked	✓ received
Routable	✓ received	✓ received
Largest spoofable neighbor prefix length	/19	/32

Meaning of result status in column:

received	Spoofed packet was received.	✓	Pattern of tests from this IP block indicates a switch from allowing spoofing to blocking it.
rewritten	Spoofed packet was received, but the source address was changed en route.		
blocked	Spoofed packet was not received, but unspoofed packet was.		
unknown	Neither spoofed nor unspoofed packet was received.		

No spoofed packets received
Private spoofed packets received
Routable spoofed packets received
IPv4
Client
7859
AS 7018
12.91.92.25
12.123.38.118
12.122.2.85
12.123.16.77
12.123.18.233
192.205.36.14
3257
6461
3356
1103
7575
38022
2152
1213
786
2914
681
195
2907
1909
7660
IPv6
Client
7859
AS 7018
2001:1890:c02:fa05::ee72:71f9
2001:1890:ff:ffff:12:123:38:118
2001:1890:ff:ffff:12:122:2:85
3356
2001:1890:ff:ffff:12:123:16:77
1213
38022
786
3257
6939
681
1103
2152
195

IPv4 Adjacent Netblock Testing:

Your host (216.146.195.x/24) can spoof 8191 neighboring addresses (within your /19 prefix)

Spoofed source address (anon)	Prefix Length	ASN of spoofed source address	Received
216.146.195.x/24	/31	7859	yes
216.146.195.x/24	/30	7859	yes
216.146.195.x/24	/29	7859	yes
216.146.195.x/24	/28	7859	yes
216.146.195.x/24	/27	7859	yes
216.146.195.x/24	/26	7859	yes
216.146.195.x/24	/25	7859	yes
216.146.195.x/24	/24	7859	yes
216.146.194.x/24	/23	7859	yes
216.146.193.x/24	/22	7859	yes
216.146.199.x/24	/21	7859	yes
216.146.203.x/24	/20	7859	yes
216.146.211.x/24	/19	7859	yes
216.146.227.x/24	/18	393713	no
216.146.131.x/24	/17	36689	no
216.146.67.x/24	/16	16399	no
216.147.195.x/24	/15	UNROUTED	no
216.144.195.x/24	/14	27553	no
216.150.195.x/24	/13	13649	no
216.154.195.x/24	/12	20141	no
216.130.195.x/24	/11	20052	no
216.178.195.x/24	/10	UNROUTED	no
216.210.195.x/24	/9	7385	no
216.18.195.x/24	/8	18450	no

Meaning of status in Received column:

no	Packets spoofed from adjacent netblock were not received, probably due to blocking
yes	Packets spoofed from adjacent netblock were received (but that netblock is within the source AS)

IPv4 Outbound Filtering Depth:

The tracefilter test found your host able to spoof routable, non-adjacent source addresses through the first 10 I

IPv6 probes:

Spoofed source address	Destination	Received
2001:48d0:101:501::159	2001:48d0:101:501::242	no
2001:4978:1fb:6400::d2	2001:48d0:101:501::242	no
2001:49aa:111:aa00::11	2001:48d0:101:501::242	no
fd11:1111:1111::1111	2001:48d0:101:501::242	no
2001:48d0:101:501::159	2001:48d0:101:501::247	no
2001:4978:1fb:6400::d2	2001:48d0:101:501::247	no
2001:49aa:111:aa00::11	2001:48d0:101:501::247	no
fd11:1111:1111::1111	2001:48d0:101:501::247	no
2001:48d0:101:501::159	2001:610:510:115:192:42:115:98	no
2001:4978:1fb:6400::d2	2001:610:510:115:192:42:115:98	no
2001:49aa:111:aa00::11	2001:610:510:115:192:42:115:98	no
fd11:1111:1111::1111	2001:610:510:115:192:42:115:98	no
2001:48d0:101:501::159	2001:630:212:225:225:90ff:fe0c:45a6	yes
2001:4978:1fb:6400::d2	2001:630:212:225:225:90ff:fe0c:45a6	no
2001:49aa:111:aa00::11	2001:630:212:225:225:90ff:fe0c:45a6	no
fd11:1111:1111::1111	2001:630:212:225:225:90ff:fe0c:45a6	no
2001:48d0:101:501::159	2001:770:18:7:225:90ff:fe0c:acb4	yes
2001:4978:1fb:6400::d2	2001:770:18:7:225:90ff:fe0c:acb4	yes
2001:49aa:111:aa00::11	2001:770:18:7:225:90ff:fe0c:acb4	yes
fd11:1111:1111::1111	2001:770:18:7:225:90ff:fe0c:acb4	yes
2001:48d0:101:501::159	2001:df0:4:800:21c:c0ff:feb2:ad5f	yes
2001:4978:1fb:6400::d2	2001:df0:4:800:21c:c0ff:feb2:ad5f	yes
2001:49aa:111:aa00::11	2001:df0:4:800:21c:c0ff:feb2:ad5f	yes
fd11:1111:1111::1111	2001:df0:4:800:21c:c0ff:feb2:ad5f	yes

IPv6 Adjacent Netblock Testing:

Spoofed source address (anon)	Prefix Length	ASN of spoofed source address	Received
2607:f441:xx::/40	/120	7859	yes
2607:f441:xx::/40	/112	7859	yes
2607:f441:xx::/40	/104	7859	yes
2607:f441:xx::/40	/96	7859	yes
2607:f441:xx::/40	/88	7859	yes
2607:f441:xx::/40	/80	7859	yes
2607:f441:xx::/40	/72	7859	yes
2607:f441:xx::/40	/64	7859	yes
2607:f441:xx::/40	/56	7859	yes
2607:f441:xx::/40	/48	7859	yes
2607:f441:xx::/40	/40	7859	yes
2607:f441:80xx::/40	/32	7859	yes
2607:f4c1:xx::/40	/24	UNROUTED	no
2607:7441:xx::/40	/16	UNROUTED	no

Meaning of status in Received column:

no	Packets spoofed from adjacent netblock were not received, probably due to blocking
yes	Packets spoofed from adjacent netblock were received (but that netblock is within the source AS)

Figure 5: *Each node in the graph corresponds to an autonomous system (AS), i.e. independent network service provider. The graph shows the path taken by spoofed packets received by our test infrastructure. IP path details are below the graph. The graph allows interactive exploration of the path to the destination, to support troubleshooting and remediation.* `https://spoofer.caida.org/report.php?sessionid=926383`

E Analysis of Remediation Efforts

To assist network operators with remediation, we created a system where they can register to receive private notifications to vetted contacts for that network. To incentivize remediation, we published reports of networks with observed SAV issues to regional network operator group email lists. We also report on *remediated* networks, i.e., networks that have deployed SAV after we found they permitted spoofed packets to exit their network.

E.1 Impact of private notifications on remediation

In initial work reported to DHS March 31st 2017 [31], we found that of the 563 ASes that we received a packet with a spoofed source IP address, 102 (18.1%) blocked packets from the same IPv4 address or IPv6 /64 network prefix in a subsequent test. While our overall remediation rate of 18.1% is encouraging, remediation is more successful in the 18 countries classified by the UK government as "majority native English speaking". Specifically, 19.7% (1 out of 5) ASes were able to provide evidence of remediation where a test that showed ability to spoof was conducted in a native English speaking country. Only 15.5% (1 out of 6) ASes were able to provide such evidence if the spoofing test was conducted outside of those 18 countries.

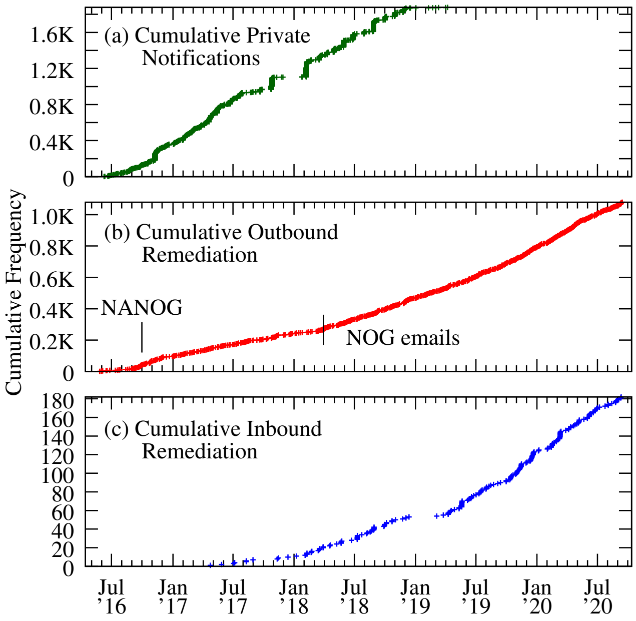

Figure 6: *Remediation events inferred between May 2016 and August 2020. Outbound remediations per month doubled when we began sending monthly NOG emails in April 2018.*

Figure 6 provides an overview of our notification and remediation activity, which began in February 2016, prior to the release of our new client system. In April 2018, we started to publicly notify members of region-specific network operator group email lists about the networks within their region that we received tests from in the past month that show gaps in SAV deployment. Figure 6a 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 6b shows the cumulative deployment of SAV for outbound spoofed packets over time; there were no corresponding bursts of remediation observed in figure 6b, leading us to believe the private notifications had limited impact.

E.2 Impact of public ("name-and-shame") notifications

We also examined the impact that our region-specific public emails to NOG email lists had on inferred remediation activity. The number of remediations we infer per month doubled from 10.6 remediations 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 6c 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.

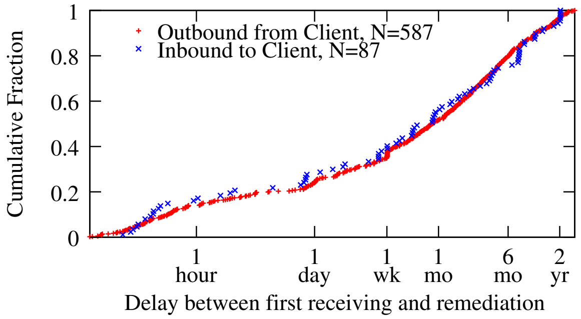

Figure 7: *Evidence of use of our system to support remediation.*

From May 2016 to August 2019, we inferred 587 instances where a network hosting a Spoofer client transitioned from forwarding spoofed packets to our system, to not doing so. Figure 7 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, 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.

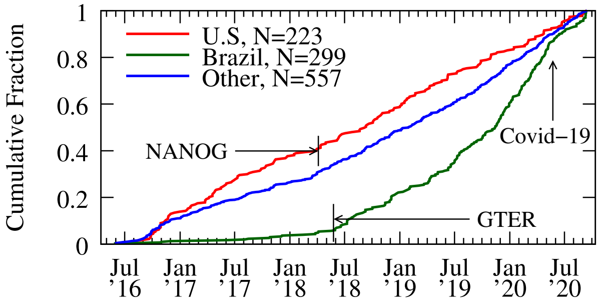

Figure 8: *Remediation in U.S. and Brazil.*

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 8 shows that nearly 90% of the remediation events in Brazil occurred after we began sending monthly emails to the Brazilian operator email list (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" [37], 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 – ≈50%. Note that remediation in Brazil has slowed since the outbreak of Covid-19 and associated shutdowns in Brazil.

F Technology Transition and Commercialization

Given the project's maturity and success, we investigated and analyzed options for transition of the technology into commercial data products. We also investigated technology transition options likely to expand SAV deployment and deliver security-relevant data into the hands of people and organizations who can ameliorate vulnerabilities. We also delivered several reports [28, 29, 35] documenting our investigation, analysis, and execution of incentive-creation scenarios that encourage SAV deployment, including feedback from a wide range of public and private sector stakeholders. We presented these results at operational (NANOG, RIPE) and academic forums (e.g., ACM Computer and Communications Security (CCS), Telecommunications Policy Research Conference).

F.1 Commercialization opportunities

We reached out to security risk management companies (such as FICO, BitSight, Security Scorecard, Shadowserver, Redseal, JSOF, GuideWire, RiskIQ), to discuss the potential for commercial use of Spoofer data or other technology transition relationships. This step precedes direct engagement with cyber-insurance companies to learn more about market requirements for such data. Without exception the feedback we received from industry was that, as with much systemic cybersecurity functionality, this technology was not sufficiently compatible with market demand to justify investment without regulatory incentives. We provide a deeper analysis of these issues in Section G.

F.2 Expansion to other measurement platforms

We facilitated transition of the software into home router (OpenWRT-focused) software devices, and platforms that use them. In particular, since 2009, the U.S. Federal Communications Commission has funded SamKnows, a U.K. company, to support the Measure Broadband America (MBA) platform. This platform provides an opportunity to expand the U.S. government's visibility of SAV compliance to many U.S. residential broadband subnets. The MBA platform is OpenWRT-based, and they provided us a test device for development, so we have verified that our our software client already runs correctly on the platform. The network load from our client is minimal, roughly 1K packets (traceroutes, spoofed probes) over the course of 5 minutes. We submitted a proposal through the formal process, which is still being evaluated by the FCC and the MBA provider.

G Overcoming misaligned incentives to deploy SAV

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. 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.

G.1 Impact of exogenous interventions

As shown in Section E.2, 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 an example from the private-sector led effort by the Internet Society: the MANRS project (Section B). Other examples are covered in our CCS2019 paper [29].

As of August 2019, the Internet Society had 205 distinct organizations (some with multiple ASes) participating in MANRS , 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 2). 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.

We reevaluated this situation in August 2020, toward the end of the project, and found some measure of improvement. At that point, the Internet Society had more than doubled the number of distinct participating organizations to 467 (some with multiple ASes) in MANRS, but only 60.3% (282) asserting their commitment to SAV on the MANRS website. For this year we had IPv4 tests from 159 MANRS ASes with no NAT, again likely the MANRS ISP member testing their own network, with only 22 (13.8%) able to spoof. We also had IPv4 tests from 209 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 (13.3%) was better than the general population (22.4%, figure 1). Although it appears that a higher fraction of those networks committing to SAV have deployed it, a much lower fraction are even committing to comply with this best practice. This is consistent with our conclusions that private-sector efforts alone will not address this vulnerability.

G.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) [12], though our measurements show that inbound SAV deployment is also problematic.

Unfortunately, there are two barriers. The first 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.

G.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 [46] and recently updated transparency requirements [47]. 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 would be an excellent arbiter of compliance with this rule, demonstrating publicly whether the network allows spoofing. This data would be useful to insurers, regulators, and to consumers wishing to understand network hygiene.

G.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 [40]. 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 [36, 42]. Sometimes NIST takes these guidance documents and embodies them in Federal Information Security Modernization Act (FISMA) controls, e.g., for Domain Name System Security (DNSSEC) [49] or in other policy initiatives [50]. 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 [45]. 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 [22]. 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 [13]. 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 [14].

Second, many people tend to look at security as the responsibility of hardware and software manufacturers. In the case of the Mirai botnet [1], the U.S. Federal Trade Commission (FTC) sued the device manufacturer (D-Link) for failing to adequately secure the company's home networking hardware [48]. We also note that a judge subsequently dismissed the lawsuit for failing to show sufficient harm by D-Link devices on consumers [44]. 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?

G.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 [20]. 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

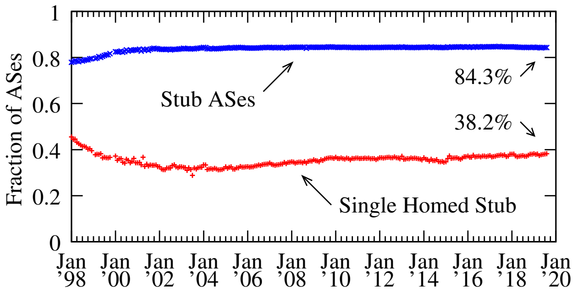

Figure 9: Feasibility of uRPF over time based on observed BGP announcements across 21 years. As of August 2019, 84.3% of ASes in the Internet are Stub ASes, and 45.3% of these stub ASes in the Internet had a single inferred transit provider (38.2% of all ASes) and were candidates for feasible-mode uRPF. (Transit relationships inferred from BGP data from RouteViews and RIPE RIS using [30].)

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) [2], 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 recently proposed improvements to filtering techniques to increase their operational robustness in the face of such complexity [41].

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 9 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 [25]. 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.

H Deliverables

YEAR 1

#	Deliverable	Type
1	Status report on deployment of anti-spoofing BP	Report [29]
2	Updated web site	Software [11]
3	Client and server software release (1st release)	Software [10]
4	Updated approach to measuring deployment of anti-spoofing BP	Report [28, 29, 35]

YEAR 2

#	Deliverable	Type
1	Status report on deployment of anti-spoofing BP	Report [29]
2	Client/server software (2nd release)	Software [10]
3	AS-level registration system	Software [9]
4	Client/server software (final release)	Software [10]
5	Improved AS-level graph visualization	Software [Figure 5]
Example: `https://spoofer.caida.org/report.php?sessionid=926383`		
6	Paper	Paper [28, 29, 35]
7	Updated approach to measuring deployment of anti-spoofing BP	Report [28, 29, 35]

References

[1] 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. Understanding the mirai botnet. In *26th USENIX Security Symposium*, pages 1093–1110, Vancouver, BC, 2017.

[2] F. Baker and P. Savola. Ingress filtering for multihomed networks. RFC 3704, March 2004. BCP 84.

[3] CAIDA. [AusNOG] spoofer report for AusNOG for Apr 2018, May 2018. `http://lists.ausnog.net/pipermail/ausnog/2018-May/040951.html`.

[4] CAIDA. [NLNOG] spoofer report for NLNOG for Mar 2018, April 2018. `http://mailman.nlnog.net/pipermail/nlnog/2018-April/002703.html`.

[5] CAIDA. [nznog] spoofer report for NZNOG for Apr 2018, May 2018. https://list.waikato.ac.nz/pipermail/nznog/2018-May/022783.html.

[6] CAIDA. Relatório spoofer para gter - Mai/2018, June 2018. https://eng.registro.br/pipermail/gter/2018-June/074470.html.

[7] CAIDA. Spoofer report for NANOG for Mar 2018, April 2018. https://mailman.nanog.org/pipermail/nanog/2018-April/094945.html.

[8] CAIDA. [uknof] spoofer report for UKNOF for Apr 2018, May 2018. https://lists.uknof.org.uk/cgi-bin/mailman/private/uknof/2018-May/005997.html.

[9] CAIDA. As spoofing notification registration page, September 2020. https://spoofer.caida.org/register.php.

[10] CAIDA. Download spoofer client software, September 2020. https://www.caida.org/projects/spoofer/#software.

[11] CAIDA. Spoofer home page, September 2020. https://spoofer.caida.org/.

[12] Alan Calder and Geraint Williams. *PCI DSS: A Pocket Guide*. It Governance Ltd, 4th edition, 2015.

[13] Charlie Osborne. 2016. https://www.zdnet.com/article/krebs-on-security-booted-off-akamai-network-after-ddos-attack-proves-pr

[14] Yi-Ching Chiu, Brandon Schlinker, Abhishek Balaji Radhakrishnan, Ethan Katz-Bassett, and Ramesh Govindan. Are we one hop away from a better Internet? In *IMC*, October 2015.

[15] Z. Duan, X. Yuan, and J. Chandrashekar. Constructing inter-domain packet filters to control IP spoofing based on BGP updates. In *INFOCOM*, 2006.

[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. The matter of Heartbleed. In *IMC*, 2014.

[17] W. Eddy. TCP SYN flooding attacks and common mitigations. RFC 4987, August 2007.

[18] P. Ferguson and D. Senie. Network ingress filtering: Defeating denial of service attacks which employ IP source address spoofing, May 2000. IETF BCP38, RFC 2827.

[19] Internet Society. Mutually Agreed Norms for Routing Security (MANRS), 2019. https://www.manrs.org/.

[20] Eric J Johnson and Daniel Goldstein. Do defaults save lives? *Science*, 302(5649):1338–1339, November 2003.

[21] Mattijs Jonker, Alistair King, Johannes Krupp, Christian Rossow, Anna Sperotto, and Alberto Dainotti. Millions of targets under attack: a macroscopic characterization of the DoS ecosystem. In *IMC*, November 2017.

[22] K. Claffy and D. Clark. Workshop of Internet Economics 2019 Final Report, 2019. https://www.caida.org/outreach/workshops/wie/.

[23] Sam Kottler. Github engineering ddos incident report, March 2018. https://githubengineering.com/ddos-incident-report/.

[24] Megan Kruse. CAIDA spoofer project improves routing security by publicizing spoofed source address packets, May 2018. https://www.manrs.org/2018/05/.

[25] Franziska Lichtblau, Florian Streibelt, Thorben Krüger, Philipp Richter, and Anja Feldmann. Detection, classification, and analysis of inter-domain traffic with spoofed source ip addresses. In *IMC*, 2017.

[26] B. Liu, J. Bi, and A. V. Vasilakos. Toward incentivizing anti-spoofing deployment. *IEEE Transactions on Information Forensics and Security*, 9(3):436–450, March 2014.

[27] Xin Liu, Ang Li, Xiaowei Yang, and David Wetherall. Passport: Secure and adoptable source authentication. In *NSDI*. USENIX, 2008.

[28] Q. Lone, M. Luckie, M. Korczyński, H. Asghari, M. Javed, and M. van Eeten. Using Crowdsourcing Marketplaces for Network Measurements: The Case of Spoofer. In *Network Traffic Measurement and Analysis Conference (TMA)*, Jun 2018.

[29] M. Luckie, R. Beverly, R. Koga, K. Keys, J. Kroll, and k. claffy. Network Hygiene, Incentives, and Regulation: Deployment of Source Address Validation in the Internet. In *ACM Computer and Communications Security (CCS)*, Nov 2019.

[30] Matthew Luckie, Bradley Huffaker, Amogh Dhamdhere, Vasilieos Giotsas, and kclaffy. AS relationships, customer cones, and validation. In *IMC*, pages 243–256, Oct 2013.

[31] Matthew Luckie and kc claffy. Strategies for region-specific SAV focus, March 2017.

[32] Sean Lyngaas. Someone is spoofing big bank IP addresses – possibly to embarrass security vendors, April 2019. `https://www.cyberscoop.com/spoofed-bank-ip-address-greynoise-andrew-morris-bank-of-america/`.

[33] Doug Montgomery and Kotikalapudi Sriram. Evaluation and Deployment of Advanced DDoS Mitigation Techniques, 2017. `https://www.nist.gov/sites/default/files/documents/2017/09/22/ddosd-nist-2016-08-v3.pdf`.

[34] Christopher Morrow. BLS FastAccess internal tech needed, 2006. `http://www.merit.edu/mail.archives/nanog/2006-01/msg00220.html`.

[35] L. Müller, M. Luckie, B. Huffaker, k. claffy, and M. Barcellos. Challenges in Inferring Spoofed Traffic at IXPs. In *ACM SIGCOMM Conference on emerging Networking EXperiments and Technologies (CoNEXT)*, Dec 2019.

[36] National Institutes of Standards and Technology. NIST Special Publication 800-series General Information, 1990. `https://www.nist.gov/itl/nist-special-publication-800-series-general-information`.

[37] NIC.br. Program for a safer internet, 2018. `https://bcp.nic.br/i+seg/`.

[38] PeeringDB. PeeringDB. `https://www.peeringdb.com/`.

[39] Zhiyun Qian, Z. Morley Mao, Yinglian Xie, and Fang Yu. Investigation of triangular spamming: A stealthy and efficient spamming technique. In *IEEE Security and Privacy*, May 2010.

[40] Tony Scott. M-15-13: Policy to require secure connections across federal websites and web services, 2015. `https://https.cio.gov/`.

[41] K. Sriram, D. Montgomery, and J. Haas. Enhanced feasible-path unicast reverse path filtering, 2019. `https://tools.ietf.org/html/draft-ietf-opsec-urpf-improvements-03`.

[42] Kotikalapudi Sriram and Doug Montgomery. Secure Interdomain Traffic Exchange - BGP Robustness and DDoS Mitigation. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP) 800-189, 2018. `https://csrc.nist.gov/publications/detail/sp/800-189/draft`.

[43] US-CERT. Multiple DNS implementations vulnerable to cache poisoning VU#800113, 2008.

[44] U.S. District Court of Northern District of California. Order Re: Motion to Dismiss: Federal Trade Commission v. D-link Systems, 2017. `https://consumerist.com/`

consumermediallc.files.wordpress.com/2017/09/dlinkdismissal.pdf.

[45] U.S. Federal Communication Commission. U.S. Anti-Bot Code of Conduct (ABCs) for Internet Service Providers (ISPs), March 2012. https://transition.fcc.gov/bureaus/pshs/advisory/csric3/CSRIC-III-WG7-Final-Report.pdf.

[46] U.S. Federal Communication Commission. In the Matter of Protecting and Promoting the Open Internet: Report and Order on Remand, Declaratory Ruling, and Order, 2015. 30 FCC Rcd 5601 (7;) 80 FR 19737.

[47] U.S. Federal Communication Commission. Restoring Internet Freedom, 2018. 33 FCC Rcd 311 (1).

[48] U.S. Federal Trade Commission. FTC Charges D-Link Put Consumers' Privacy at Risk Due to the Inadequate Security of Its Computer Routers and Cameras, January 2017. https://www.ftc.gov/news-events/press-releases/2017/01/ftc-charges-d-link-put-consumers-privacy-risk-due-inadequate.

[49] U.S. Government Office of Management and Budget. Securing the Federal Government's Domain Name System Infrastructure, 2008. https://www.whitehouse.gov/sites/whitehouse.gov/files/omb/memoranda/2008/m08-23.pdf.

[50] U.S. Government Office of Management and Budget. Internet Protocol Version 6 (IPv6), mandated transition, 2010. https://obamawhitehouse.archives.gov/sites/default/files/omb/assets/egov_docs/transition-to-ipv6.pdf.

[51] A. Yaar, A. Perrig, and D. Song. StackPi: New packet marking and filtering mechanisms for DDoS and IP spoofing defense. *IEEE JSAC*, 24(10), Oct 2006.