# Technical Proposal - Augment Spoofer Project to Improve Remediation Efforts (ASPIRE)

Dr. Kimberly Claffy

# Table of Contents

# A Abstract

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 is to promote using SAV BCP by networks around the world.

With previous DHS funding we have re-designed, re-implemented, deployed, and operated a secure measurement infrastructure, Spoofer, that supports large-scale studies of anti-spoofing measures deployed (or not) in the global Internet. In the process, we have demonstrated the soundness of the technical concepts of our measurement approach and implementation, and established valuable relationships with the operational security community who have provided us feedback and encouragement in expanding and publicizing the project and data. However, during the course of the project we have realized that there is a gap between generating security hygiene data and achieving remediation at scale. Thus, after successful completion of all tasks funded by the contract D15PC00188 (Software Systems for Surveying Spoofing Susceptibility), we propose the following new tasks targeting focused remediation efforts, for a new 2-year contract, in an international collaboration with the University of Waikato.

Our proposed tasks include: maintaining the Spoofer client-server platform operations and improving the project reporting web site to facilitate remediation efforts; investigating, evaluating, pursuing, and documenting the effects of different approaches to stimulating remediation activities, including integration of data into security risk management and commercial cyber-insurance ecosystem; and analyzing, socializing, and documenting community feedback on economic and regulatory options to support SAV deployment. These tasks aim at achieving greater involvement from a broader cross-section of security research, operations, risk management, and public policy stakeholders.

We are uniquely qualified to pursue this work. First, we have extensive experience obtained from developing and operating the current Spoofer project that leads toward the new developments we propose in this document. Second, we control and operate unique measurement infrastructures: an Internet-scale active measurement platform, which helps us to assess and report the status of SAV deployment, and the UCSD network telescope, which we can use to corroborate our conclusions regarding the observable effects of SAV policy on spoofed DDoS attacks prevalence. Finally, we have trust and respect of the Internet operational community that lends greater weight to our remediation efforts.

# B Performance Goals

The Regents of the University of California; University of California, San Diego on the behalf of the San Diego Supercomputer Center's Center for Applied Internet Data Analysis (CAIDA) research program, offer this technical proposal which includes the following deliverables: (1) updates to the production-quality client-server source address validation (SAV) testing system that we built in the previous contract to further scrutinize results obscured by network address translation; (2) a reporting and analysis system that stimulates remediation activities through geographic-focused operator notifications and enables assessment of their impact; (3) an Autonomous System (AS) level registration system that allows network operators to sign-up for notifications for when we receive tests that show lack of source address validation; and (4) a report that documents executed and proposed approaches for integration of data into the security risk management and commercial cyber-insurance ecosystem.

The project will leverage the results of existing technologies and infrastructure funded by the Department of Homeland Security and the National Science Foundation.

The proposed work targets objectives outlined in TTA#1: Measurement and Analysis to Promote Best Current Practices. Specifically, we propose to refine and operate multiple open-source software tools for anti-spoofing assessment that will allow a site to determine if it has successfully deployed source address validation, and provide on-going monitoring and testing to ensure SAV continues to operate correctly through network upgrades and reconfigurations. Our reporting and analysis system will promote the deployment of SAV by guiding compliance attention where it will have the most benefit, and provide independent measures of the effectiveness of our approaches to promoting SAV best practices. To enable additional testing that will magnify our view of SAV deployment on many networks, we will pursue three additional goals: (1) testing of networks where a Network Address Translation device rewrites the source address of packets with spoofed source addresses, obscuring our visibility into whether that device properly blocks spoofed packets; (2) expanded public notifications and reporting through our operator-focused reporting engine; (3) development of incentive-creation scenarios, e.g., economic (cyber-insurance ecosystem) and policy approaches that encourage remediation.

The resulting technologies and data will improve our ability to identify, monitor, and mitigate the infrastructure vulnerability that serves as the primary vector of massive DDoS attacks on the Internet.

# C Detailed Technical Approach

Despite source IP address spoofing being a known vulnerability for close to 30 years [2], and despite many efforts to mitigate or even shed light on the problem (e.g. [3, 4, 5]), spoofing remains a viable attack method for redirection, amplification, and anonymity, as evidenced most recently and publicly in March 2018 during a 1.7 Tbps DDoS attack against Github [16]. That particular attack used an amplification vector in Memcached [16]; previous attacks against Cloudflare [18] and Spamhaus [6] in 2013 achieved 300+ Gbps using amplification vectors in NTP and DNS. In all of these cases, the attacks exploited the ability of (many) publicly accessible networks to spoof IP packets. While some application-layer patches can mitigate these vulnerabilities [19], attackers continuously search for new vectors. To defeat spoofed-source DDoS attacks requires operators

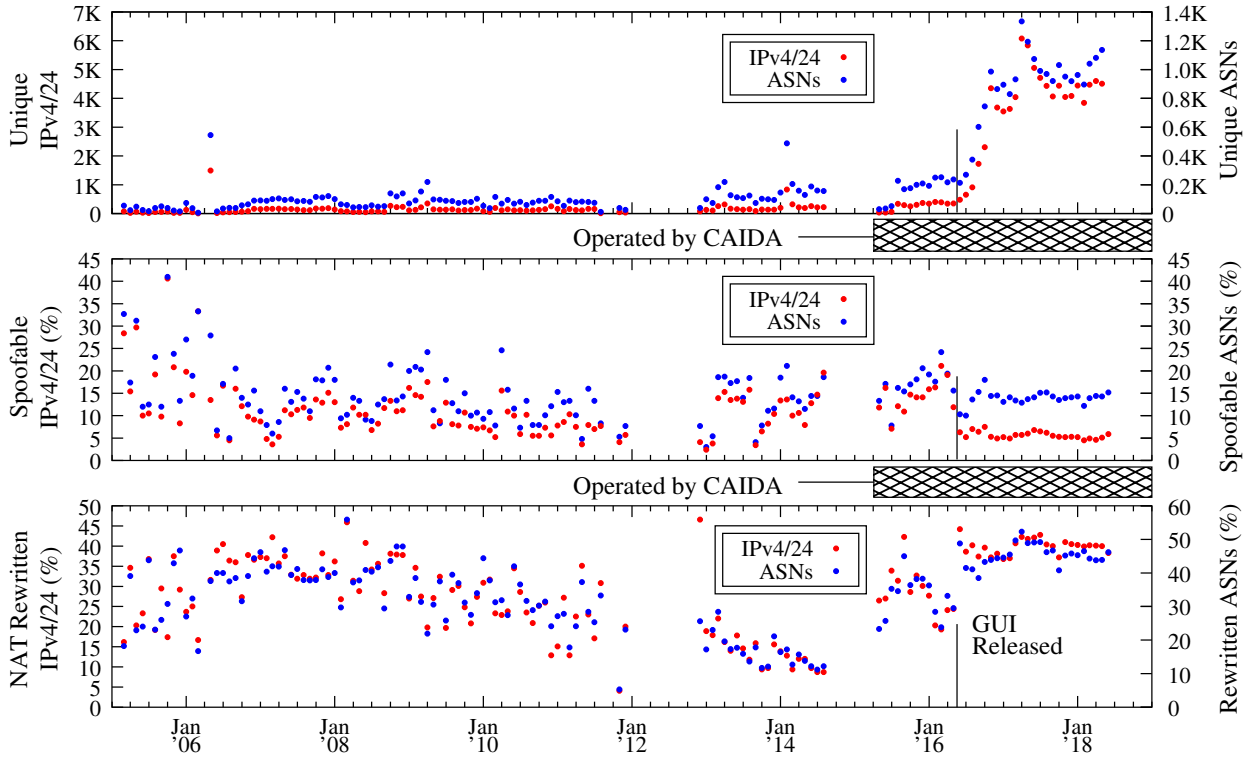


Figure 1: *Overview of Spoofer project data collection over time, aggregated per month. The gaps prior to May 2015 are due to hardware failures. After we released our new software system in May 2016, the volume of tests have increased an order of magnitude from ≈300 IPv4 /24 prefixes in ≈200 ASes to ≈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 represent a gap in our visibility into SAV deployment. We propose a new measurement technique (Figure 3) to close this gap.*

to ensure their networks filter packets with spoofed source IP addresses [13], 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.

Testing a network's SAV compliance requires a measurement vantage point inside (or immediately upstream of) that network, because the origin network of arbitrary spoofed packets cannot be determined [1]. During the past three years, our approach was to build a production-quality software client that volunteers across the Internet could download and run from their networks, testing their own network's ability to send various types of spoofed packets to our server, which collected, aggregated, and publicly reported test results. Our current 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 where along a path SAV may be tak-

ing place. A key improvement we made to the previous software client was to run spoofing tests periodically in the background, initiating tests on any attached networks once per week, which allows us to study longitudinal SAV deployment.

We have used the resulting data to inform the continuing debate in the operational security community on which networks on the Internet permit spoofed packets to exit their networks. We have also publicly reported anonymized test results for ASes where we have received tests, contacted networks with outcomes of tests conducted from their network where we received packets with spoofed source IP addresses, publicly reported networks that have deployed SAV after we found they permitted spoofed packets to exit their network, and started to publicly report networks with apparent SAV issues to region-focused network operator group email lists. Figure 1 provides an overview of the Spoofer Project's data collection since the project began in 2005. After we released our new software system in May 2016, 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 (top panel of figure 1). We found that 15% of these tested ASes have not deployed SAV uniformly throughout their network (middle panel of figure 1).

Figure 2 provides an overview of our notification and remediation activity. Our notification activity commenced in February 2016 (prior to the release of our new client system). While some networks deploy SAV without our notifying them that we received a positive Spoofer test from that network, figure 2 shows that bursts of remediation activity are correlated with bursts of private notification from us. Beginning 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. Currently, we notify operator lists covering networks in the US and Canada, the Netherlands, the United Kingdom, Australia, New Zealand, and Brazil on the 8th day of each month. We have noticed bursts of remediation activity correlated with these notifications in the countries covered by them.

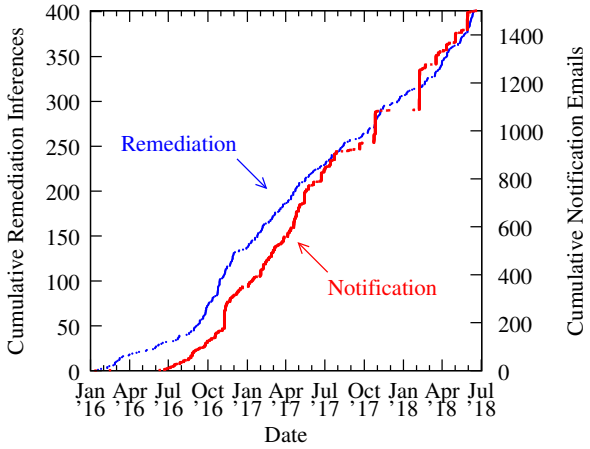


Figure 2: *Correlating remediation with notification. Remediation occurs at a lower rate during periods where we did not send private notifications than during periods when we did.*

With our success in obtaining and reporting data on SAV deployment, and effecting remediation, we seek support to expand our view of SAV deployment and increase remediation. We propose three complementary efforts: improving capabilities of the Spoofer software to close visibility gaps with new measurement techniques, *methods to improve remediation outcomes, including integration of data products and capabilities into the commercial ecosystem supporting security risk management;* and maintaining the production-quality client/server system to ensure continued deployment.

Our initial development focus will be in improving the capabilities of the Spoofer software to support testing of SAV when an intermediate router along the tested path rewrites the source address of spoofed packets. Our current system classifies a test where a Network Address Translation

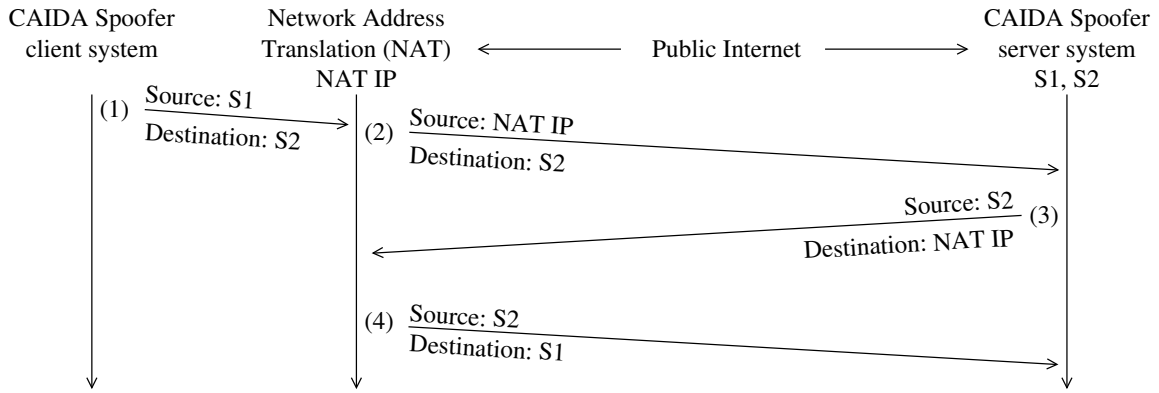


Figure 3: *Extending the Spoofer client-server system to test clients whose packets with spoofed source addresses are rewritten. Our current system sends packets (1) and (2), and classifies the client's spoofed packets as rewritten because source address S1 in packet (1) was changed to the NAT router's 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 from the NAT IP to S1, even though S1 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.*

(NAT) router rewrites the spoofed source address as *rewritten*. Currently, we observe this behavior for 37.2% – 42.8% of IPv4/24 prefixes tested per month (bottom panel of figure 1) – representing not only a visibility gap, but an untapped set of vantage points for testing SAV deployment, with six times more prefixes than we receive spoofed packets from. When the NAT router rewrites the source address of a packet, it creates an internal mapping, so that the NAT router can forward any response it receives to the appropriate system. We have 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, as illustrated in figure 3. In figure 3, 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.

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 [17]. 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 will create a registration system in the Spoofer project's reporting engine that allows a vetted operator within an AS to receive these notifications.

Further, representatives in the operational security community encouraged us to begin send-

ing public notifications where we infer SAV is not deployed. Beginning March 2018, we have sent monthly emails to six public region-focused network operator group (NOG) emailing lists – NANOG covering the United States and Canada [11], NLNOG covering the Netherlands [8], AusNOG covering Australia [7], NZNOG covering New Zealand [9], UKNOF covering the United Kingdom [12], and GTER covering Brazil (which we translated to Portuguese [10]). 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 have received public support for this activity from the Internet Society (ISOC) Mutually Agreed Norms for Routing Security (MANRS) initiative [14]. Therefore, we will expand our public notifications to cover more NOGs, translated into the native language spoken by members of the NOG. (A German NOG recently requested us to add their region to our monthly reporting service.)

Next, we will examine the impact our notification activities have on remediation activities. In initial work reported to DHS March 31st 2017 [15], 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. However, 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.

We will examine the impact that our region-specific emails to NOG email lists have on inferred remediation activity. Given the project's maturity and success thus far, we will also investigate and analyze options for transition of the technology into commercial data products. We will begin this task by reaching out to security risk management companies, e.g., FICO, BitSight, Security Scorecard, Shadowserver, and Redseal, to discuss the potential for commercial use of Spoofer data or other technology transition relationships. This step will precede direct engagement with cyber-insurance companies to learn more about market requirements for such data. We will also facilitate transition of the software into home router (OpenWRT-focused) software platforms, which may lead to commercial licensing of the software for use in other home router platforms. We will document the results and analysis of these and other scenarios to incentivize deployment.

Finally, we will continue to maintain and operate the existing infrastructure, ensuring that our software continues to function on modern operating systems as updates to operating systems occur. In particular, we will update the operating system platform that our server infrastructure uses to ensure the platform is supported by the operating system maintainers. We will also continue to test our client software on Windows, MacOS, and Linux to ensure that our software continues to work. In December 2017, we had to modify our client software to prevent the network stack in MacOS from becoming unstable, due to a bug in MacOS that Apple introduced with the High Sierra release.

## C.1 Improve capabilities of Spoofer software

To allow for improved visibility into Internet Service Provider (ISP) deployment of SAV, we will extend the capabilities of our Spoofer software to further test networks whose Network Address

Translation (NAT) router obscures our view of SAV deployment. We will modify our server software to respond to packets whose spoofed source address has been rewritten in order to infer SAV policy for these networks. Importantly, these changes only require modification of the server software, as our modifications only interact with the NAT router in question, implying that all currently deployed client software will benefit from this work.

### C.2 Explore methods to stimulate remediation

We will extend our reporting engine to contact further region-specific network operator group email lists. The Wikipedia page lists ≈50 country-specific network operator groups, of which we currently notify six. For the NOGs that do not contain majority native English speakers, we will seek translations to ensure our remediation emails are understandable.

We will also create a web-based registration system that allows networks to self-register vetted contacts within their AS. 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 will send the email to the registered contact.

We will investigate and evaluate technology transition options that are likely to expand SAV deployment and deliver security-relevant data into the hands of people and organizations who can ameliorate vulnerabilities. We will contact security risk analysis companies (such as FICO, BitSight, Security Scorecard, Shadowserver, and Redseal), to discuss the potential for use of Spoofer data, and integration of Spoofer data into commercial products they sell to insurance companies. We will document the results of this research to inform our analysis of scenarios to incentivize deployment, including policy and regulation options, and community feedback on various scenarios that encourage remediation for networks that will not otherwise deploy SAV. We will engage with relevant government agencies (NIST, FCC, DHS, DoC) regarding their view of their role in promotion or enforcement of SAV deployment, and integrate their feedback into the report. We will socialize these results at CAIDA workshops, operational meetings, e.g., NANOG, RIPE, and policy research forums, e.g., Telecommunications Policy Research Conference (tprcweb.com). If there is sufficient interest, we will maintain and update this document as SAV-related policies evolve.

### C.3 Maintain Spoofer operations

We will continue 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 will upgrade the operating system on the deployed spoofer servers to ensure the underlying software is still supported by the vendor. We will expand our efforts to build Spoofer packages for home access router platforms, e.g., OpenWRT-based, and investigate options for integration of the software into other home access router platforms.

## D Testing and Evaluation

CAIDA has ready access to computer systems running the operating systems required to test and develop our client and server SAV testing software (MacOS X, Linux, Windows). CAIDA also op-

erates the Archipelago measurement infrastructure that allows our client-server system to evaluate of the placement of SAV filters along Internet paths. We will utilize local resources and expertise to build, test, and evaluate all software deliverables. We will use open-source static analysis systems (e.g., Clang Static Analyzer and cppcheck) and dynamic analysis systems (e.g. Valgrind and dmalloc) to audit our tools as we develop them, and then utilize the capabilities of the DHS Software Assurance Marketplace (SWAMP) to audit our completed client-server system. We will solicit feedback from DHS and network operators on new versions of our systems as we build them to ensure our software is designed and implemented to have the highest utility to all stakeholders.

# References


[1] F. Baker and P. Savola. Ingress filtering for multihomed networks, March 2004. IETF BCP84, RFC 3704.

[2] S.M. Bellovin. Security problems in the TCP/IP protocol suite. *ACM/SIGCOMM Computer Communication Review (CCR)*, 19(2):32–48, April 1989.

[3] Robert Beverly and Steven Bauer. The spoofer project: Inferring the extent of source address filtering on the Internet. In *Proceedings of USENIX SRUTI*, July 2005.

[4] Robert Beverly, Arthur Berger, Young Hyun, and k claffy. Understanding the efficacy of deployed Internet source address validation filtering. In *Proceedings of the 9th ACM SIGCOMM Internet Measurement Conference*, November 2009.

[5] Robert Beverly, Ryan Koga, and kc claffy. Initial longitudinal analysis of IP source spoofing capability on the Internet, July 2013. `http://www.internetsociety.org/`.

[6] Peter Bright. Spamhaus DDoS grows to Internet-threatening size, March 2013.

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

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

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

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

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

[12] 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`.

[13] 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.

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

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

[16] Lily Hay Newman. A 1.3-Tbs DDoS hit GitHub, the largest yet recorded, March 2017. `https://www.wired.com/story/github-ddos-memcached/`.

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

[18] Matthew Prince. Technical details behind a 400Gbps NTP amplification DDoS attack, February 2014. https://blog.cloudflare.com/technical-details-behind-a-400gbps-ntp-amplification-ddos-attack.

[19] Paul Vixie. Rate-limiting state: The edge of the Internet is an unruly place. *ACM Queue*, 12(2):1–5, February 2014.