SAVing the Internet: Explaining the Adoption of Source Address Validation by Internet Service Providers

Qasim Lone[1], Maciej Korczyński[2], Carlos H. Gañán[1], and Michel van Eeten[1]

[1] Delft University of Technology, The Netherlands
[2] Univ. Grenoble Alpes, CNRS, Grenoble INP, LIG, F-38000 Grenoble, France

Spoofed Internet traffic is used by miscreants, most visibly for amplification DDoS attacks. Source Address Validation (SAV) by network operators is a security best practice to stop spoofed traffic from leaving their network. Its adoption is hampered by incentive misalignment: the cost is borne by the operator, while the benefits go to the rest of the Internet. This paper estimates the impact of various incentives on SAV adoption. It is the first study that combines two independent datasets with observations for the absence of SAV and that statistically models its causal drivers. We map these observations to a population of 334 ISPs that control the bulk of the market share for Internet access in 61 countries. We find evidence for the absence of SAV for certain prefixes of 250 ISPs. Next, we try to explain what portion of an ISP's address space allows spoofing from four causal factors – network complexity, security effort, ISP characteristics and institutional environment – as measured via 12 indicators. We find evidence larger ISPs have a higher proportion of non-compliant IP space. ISP security efforts, most notably the adoption of RPKI and the number of amplifiers, are positively related to SAV. Subscription prices and ISP revenue have no significant impact. Finally, we find that ISPs in countries with more developed ICT infrastructures are also more likely to have a wider adoption of SAV. We reflect on these findings and discuss potential ways forward for SAV.

1 Introduction

Spoofed Internet traffic—crudely put, IP packets with forged source IP addresses—has been a persistent security problem for decades. Used in a variety of attacker practices, its most visible consequence has been the problem of Distributed Denial-of-Service (DDoS) attacks based on amplification. This has led Internet Hall of Fame technologist Paul Vixie to conclude: 'Nowhere in the basic architecture of the Internet is there a more hideous flaw than in the lack of enforcement of simple SAV (source-address validation) by most gateways' [47].

The way attackers abuse spoofed traffic in amplification DDoS attacks is by sending queries with a forged source IP to a multitude of servers running amplification protocols, such as open DNS resolvers or Memcached servers. The spoofed packets set the victim's IP address as the source IP address. As a result, the victim receives a large number of server responses that congest its network or system, making it unavailable for incoming and outgoing traffic. This attack is

hard to mitigate by the victim's network. However, if network operators would verify the source address of the packets originating from their own networks, and drop illegitimate packets, it would curtail the ability of the attacker to successfully send spoof packets in the first place. This practice is commonly known as Source Address Validation (SAV), most notably documented in BCP38. The idea of BCP38 is that a network operator checks the source address of every outgoing packet before it leaves its network against a set of allocated addresses. They should drop the packet if the source address is outside the range of IPs assigned to them.

Even though for years now there has been a push for the implementation of BCP38 across operators, we still observe 22% of all observed ASes being non-compliant in the Spoofer dataset [21]. The adoption of BCP38 suffers from a clear misalignment of incentives: the cost is borne by the network that adopts it, while the benefits go to the rest of the Internet. Non-compliance can therefore be seen as causing a negative externality. Seen in this light, it is actually remarkable that a sizeable portion of all networks *are* in fact compliant.

This paper presents the first study that measures the current state of SAV using two independent measurement techniques and that identifies causal factors for non-compliance. We use the terms compliance and adoption interchangeably. The underlying causal mechanisms are likely to be different, if not outright incomparable, across the enormous heterogeneity of operators behind the more than 60,000 Autonomous Systems (AS) that currently make up the Internet. For this reason, we focus our analysis on a critical population with a more homogeneous composition: Internet Service Providers (ISPs), here defined as the businesses that offer Internet access to end users. Given that these networks offer access to billions of users, they are also a critical control point for adopting SAV and block potential miscreants from IP spoofing. What is also important here is that the BCP38 unambiguously applies to such so-called 'stub' networks that ISPs operate, as opposed to the more complicated case of transit networks [31]. The underlying problem for transit providers is that they might have customers that are not announcing routes to them, due to traffic engineering. If the non-stub network drops these IP packets, they are losing legitimate traffic destined towards its downstream customers. BCP84 introduced several improvements to BCP38 and proposed filtering using static Access Control Lists (ACLs) or Reverse Path Forwarding (RPF) [15]. It also suggests that in the case of asymmetric routing, network operators should only drop packets with "martian addresses" or currently not routed IP addresses. This will limit the problem, but the routable IP space can still be spoofed. Since our dataset only deals with ISPs, there are very few transit networks. Some of these transit ASes can also be siblings of stub Ases, belonging to the same organization where routing information can be shared.

The research question that we set out to answer is: What factors explain the extent in which Internet Service Providers are not compliant with BCP38? We will be looking at factors like network complexity, security effort and the institutional environment of the country where the ISP is located.

After we discuss the related work, we will unpack various economic factors that shape the incentives for compliance. We then explain how BCP38 non-compliance was measured across our study population of 250 ISPs in 61 countries. We collect various indicators for our theoretical framework and then estimate an OLS regression model to explain non-compliance. Finally, we discuss our findings in light of current industry proposals on how to increase the adoption of BCP38.

2 Background and Literature Review

2.1 SAV – Source Address Validation

Best Current Practice 38, also referred to as Source Address Validation (SAV), was proposed in RFC 2827 almost 20 years ago to respond to a growing problem of DoS attacks [41]. The RFC describes the straightforward idea of ingress filtering, which assumes that source IP addresses should be checked against a set of allowed addresses and discarded if they are not following filtering rules. If a network provider is aggregating routing announcements for its single-homed client networks, it should strictly prohibit traffic which claims to have originated from outside of these networks. RFC 3704 proposed different ways to implement ingress filtering using static Access Control Lists (ACLs) or Reverse Path Forwarding (RPF) [15].

Adversaries take advantage of the absence of SAV to launch DDoS attacks by exploiting public services vulnerable to reflection. In a typical scenario, end-user machines send requests from networks that allow spoofing to public services by forging source IP addresses of the victim [40]. The victim is then overloaded with the traffic coming from the public services rather than from the compromised machines. Therefore, the origin of the attack is not traceable.

2.2 Inferring SAV Deployment

Numerous papers proposed methods to infer SAV deployment [17, 24, 28–33, 36]. The approach of the Spoofer project [17] is to enable volunteers and "workers" remunerated through five crowdsourcing platforms in a pilot study [30] to test SAV compliance of their networks with a custom client-server system. The client sends spoofed and non-spoofed traffic to the server periodically or when it detects a different network. The server infers if SAV is deployed in a tested network. Even though the Spoofer project provides the most confident picture of the deployment of SAV, those that are unfamiliar with the problem or do not implement BCP38 are less likely to run the Spoofer client on their networks.

Another approach proposed by Mauch [33] and implemented by Kührer et al. [28] leverage the misconfiguration of DNS resolvers. DNS servers perform resolutions of human-readable domain names to IP addresses interpretable by machines. Local DNS resolvers can be configured to *forward* requests to other DNS servers that perform resolutions on their behalf. When a misbehaving *open*

DNS resolver receives a request from an external client, it may forward the request to another DNS resolver located outside its network without changing the packet source IP address to its address. If SAV is not deployed at the network edge, the client will receive the request resolution from the IP address of another resolver. The method is more practical than the Spoofer because the measurement can be performed remotely and does not require volunteers inside the tested networks.

Lone et al. [31] proposed another technique using routing loops appearing in traceroute data to infer inadequate SAV at the transit provider edge. When packets are sent to a customer network with an address that is routable but not allocated, and the default route is set to provider, the packets will be forwarded back to the provider's router without changing the packet source. If the router does not perform SAV, the traceroute will show a forwarding loop as the provider's router will again return subsequent packets to the customer's network.

Lichtblau et al. [29] and Müller et al. [36] have passively analyzed traffic at Internet Exchange Points (IXPs) to infer which source addresses should legitimately appear across IXP parts by leveraging Autonomous System (AS) topology extracted from Border Gateway Protocol (BGP) data. Even though the proposed detection method does not depend on volunteers running any custom software or existing misconfigurations, it requires privileged access to the traffic exchange points and cannot be easily replicated without special access to the data.

Luckie et. al [32] analyzed Spoofer dataset and ran remediation campaigns. They found at least a quarter of ASes did not filter packets for the year ending Aug 2019. They also found networks behind Network Address Translation (NAT) not always perform SAV. Finally, they analyzed remediations and found that 21% of networks remain unremediated for more than six months.

While the above-proposed methods infer SAV deployment for outbound traffic (i.e., coming from inside the customer network to the outside), Korczyński et al. proposed a new technique to identify networks not filtering *inbound* traffic to the customer network [22, 24, 25]. It consists of identifying open and closed DNS resolvers handling requests coming from the outside of the network with the spoofed source address from the range assigned inside the network under the test. This method covers roughly 50% of all ASes and provides the most complete picture of the status of inbound SAV deployment at network providers.

This paper is the first study to combine two independent measurement techniques (based on Spoofer and DNS Resolvers) to identify the lack of outbound SAV, as well as the first to statistically model causal factors for SAV non-compliance at the ISP level.

2.3 Modeling Security Performance

A few studies have explored concentrations of abuse events across different types of Internet intermediaries, with the intent to explain what factors correlate with abuse levels. Tajalizadehkhoob et al. [44] and Noroozian et al. [37]

explored analytical models to estimate the security performance of the hosting providers. By building generalized linear models (GLM) for phishing abuse counts, they demonstrated that hosting providers' structural properties, such as domain names space size or IP space size, but also factors reflecting security performance can predict a large amount of the variance in abuse incident counts.

Other studies have explored factors driving domain abuse of operators of Top-Level Domains (TLD) [26, 27]. They concluded that apart from structural properties of the operators, security efforts such as strict policies of domain names registration significantly reduce the number of domains used in phishing and malware attacks.

Our work is closely related to [48] in which Zhang et al. systematically explored the relationship between the mismanagement of networks using Internet-scale measurements of BGP routers, SMTP, HTTP and DNS servers, and malicious activities. They found a statistically significant correlation between networks that are mismanaged and networks that are responsible for distributing spam, malware, or phishing attacks. In this work, we collect various indicators reflecting network properties, security efforts, institutional factors and characteristics of ISPs to explain the absence of SAV using the data from the Spoofer project and measurements of misbehaving open DNS resolvers.

3 Theoretical Framework

Several economic concepts help understand the incentives of network operators to adopt or ignore best security practices like SAV. We first discuss these concepts and then present the causal framework that is the basis for our empirical study.

3.1 Incentives

Cost of adoption: First, the most obvious incentive against adoption is the demand for resources, including technical expertise, time, and hardware requirements for the implementation of SAV. The two well-known methods to deploy SAV are Access Control Lists (ACLs), which requires manually maintaining a list of all the prefixes announced by the AS, and Universal Reverse Path Filtering (uRPF), where the router checks if a source address exists in its routing table before forwarding it. Other than the requirements for implementing SAV, organizations also face ongoing maintenance costs, e.g., engineering time needed for keeping the ACL-based filtering up to date or hardware requirements for uRPF to maintain good throughput rates.

Externalities: An externality can be defined as the cost or benefit that affects a third party without this being reflected in the market price. SAV adoption suffers from externalities because the cost of adoption is borne by the operator, while the benefits go to others, e.g., the victims of amplification DDoS attacks. Simply put, operators do not see a direct economic benefit to implementing SAV in their networks. While one could argue that the cost of delivering spoofed

traffic also implies a cost to the operator where it originates [35], this effect is seen to be very small. In a survey on SAV adoption, the majority of respondents said that spoofed traffic constitutes only a small fraction of all traffic in terms of total volume [29].

Information asymmetry: Whether a network operator is compliant with SAV is often not visible to customers, other providers or outside observers. Adopting this good practice, therefore, doesn't generate a benefit in terms of a better reputation, as the information is not readily available to the public or to other providers who might use it in peering decisions. Conversely, non-compliance doesn't generate a clear negative reputation impact.

Weakest link: Finally, SAV suffers from being a weakest-link problem. If there are even a handful of non-compliant networks, the attack will remain possible. It would be difficult to trace it back to the offending network where it originates. Innovators and early adopters can definitely help the cause by reducing the number of vantage points from where an attack can be launched. However, it would not be possible to eradicate the attack vector until all of the operators are compliant. Since SAV adoption is a good practice, and there are no regulations or fines, it is unlikely all the operators will become fully compliant.

3.2 Explanatory Factors for SAV Compliance

In light of the above-mentioned incentives, we are developing a causal model that hypothesizes the cost of adoption to impact adoption. We approximate this cost in two ways. First, the more complex and dynamic the operator network is, the more costly SAV adoption will be. We include this variable as 'network complexity'. Second, if an operator has a large customer base, it will have economies of scale and be more likely to have expertise in network engineering, making it less costly to implement SAV adoption. This factor is included as 'ISP characteristics'.

The impact of the cost of adoption is moderated by other factors. First, the willingness of the operator to incur costs for security efforts. Second, by the overall development level and wealth of the country in which it operators ('institutional environment')—in other words, the extent to which they can past these costs on to their customers.

For each of these four factors, we identified several indicators that can be empirically observed. Figure 1, shows an overview of variables and the indicators to understand non-compliance of SAV. The dependent variable is defined as non-compliance because of the way compliance is measured. As we will discuss in Section 4, the two measurement techniques are able to observe the lack of compliance, rather than its complement.

Network Complexity: We hypothesize that the more complex and dynamic a network is, the more costly it will be to implement SAV safely; thus, the

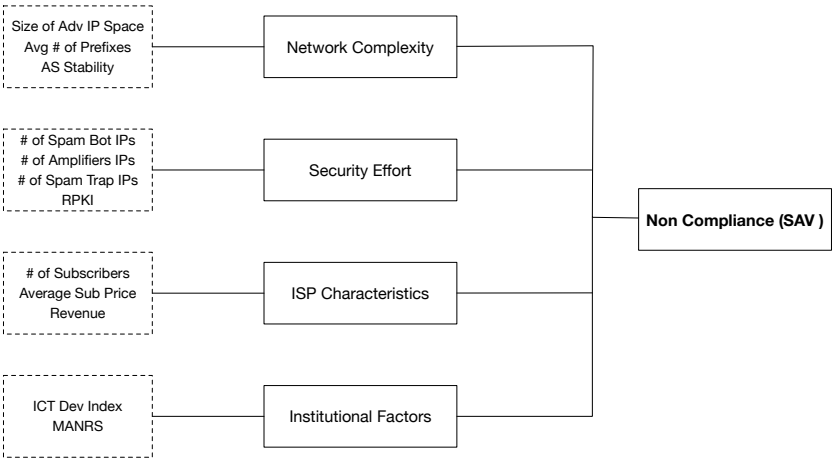

Fig. 1. Causal Model for Non-Compliance with Source Address Validation

more likely it is that the operator will not be compliant. We measure network complexity from several observable network properties. One of the important indicators is the amount of IPv4 address space advertised by ISPs. It gives us a proxy of the size of the ISP. If the operator is announcing a large number of IPs, it is more likely that they have a more complex network. They might also be running various network policies for different IP ranges, which would mean they are required to apply SAV at multiple points in the network.

Similarly, we calculated the stability of ASes based on the number of prefixes that are changing over time. The more prefixes there are, the more costly it will be to maintain the ACLs needed to implement SAV. We calculated the total number of prefixes advertised per week by ASes for April-Sept, 2019. The more the prefixes change, this will also increase the cost of adoption. We calculated the ones that remain consistent throughout the year. If advertisements are constantly changing, it would mean that it is difficult for ASes to implement ACL-based BCP38 implementation.

Security Effort: We used network hygiene to understand how well the networks are maintained. The idea behind network hygiene is to measure proxies for how much effort operators put into keeping their networks secure. We use two factors to calculate the hygiene of the networks. First, the presence of amplifiers in the network. Services like Open Resolvers, Memcached servers, and Chargen are constantly used by attackers to redirect and amplify a DDoS attack. There has been a consistent effort in the operator community to get operators to reduce the number of such amplifiers in their network. We calculated the number of amplifiers per ASN and used it as an indicator that the network operators with a higher count of amplifiers are less likely to have SAV in place.

Second, we calculated abuse in the network in the form of the number of bots and the number of spam-sending IPs. We hypothesize that the network operators who perform poorly on keeping their networks clean are less likely to care about SAV.

Finally, we checked if operators have signed one or more of their prefixes using Resource Public Key Infrastructure (RPKI). We assigned a binary value for adoption. RPKI is a framework that allows service providers to sign the prefixes allocated to them. It allows other ASNs to validate the ownership of the advertised prefix.

Since RPKI is a relatively new framework, and it is an opt-in service, we hypothesize that the operators that sign one or more prefix are security-aware and are more likely to adopt SAV for their networks.

ISP Characteristics: We define ISP characteristics as a function of the number of subscribers and available funds of ISPs. These characteristics give us a different picture compared to network properties. For instance, IPv4 allocation is not evenly distributed. It has a limited pool and was assigned based on a first-come, first-serve basis. As a result, we have many ISPs with a large customer base and fewer IP addresses.

We use subscriber numbers from Telegeography data [45] and manually map company names to ASes. Moreover, we use an average subscription price as a proxy to their earnings for available funds. We hypothesize that ISP with a large number of subscribers will face difficulties for the implementation of SAV. However, ISPs with a higher subscription price will have more funds to invest in good security practices like SAV.

Institutional Factors: Additionally, we measure the impact of institutional characteristics on ISPs for the implementation of BCP38. We expect ISPs in countries with more mature ICT development more likely to be compliant as there will be more resources, more mature networks, and more initiatives for better security. For this, we use the U.N.'s ICT Development Index [23].

Similarly, we test whether ISPs are a signatory of Mutually Agreed Norms for Routing Security regulations (MANRS). MANRS initiative recommends best practices for ISPs to reduce the most common routing threats. It requires SAV for single-homed ISPs or for those whose customer network is owned by the ISP. There are currently 209 member ISPs. We expect that these ISPs are likely to be more compliant compared to non-members.

In sum, the basic idea is that the cost and benefit of SAV adoption are highly asymmetrically distributed, making adoption much less likely. We want to estimate the impact of various causal factors on adoption. First, network complexity (via various indicators), where we assume this would increase the cost of adoption and thus lower the probability of adoption. Second, security effort, where we assume this reflects the willingness of network operators to invest in security measures that also, or even primarily, benefits third parties. We assume this increases the likelihood of adoption. Third, institutional factors,

where we assume that more operators in more wealthy countries and with more mature networks and regulatory environments are more likely to accept the cost of adoption as a 'cost of doing business', thus increasing the likelihood of adoption.

4 Data Collection

In this section, we describe the sources for various datasets that we use to estimate the model. As explained before, we focus our analysis on ISPs to have a somewhat homogeneous study population, but also because the majority of the user devices are in an ISP network. They form a critical control point because they are closer to the origin of the traffic and can not only detect but also block spoofed traffic.

Several of our datasets of the independent and dependent variables are based on IP addresses and Autonomous System Numbers (ASNs). The relationship between ASes and ISPs is not that simple. Yes, many ISPs have a single AS, but a fraction of ISPs have multiple ASNs, some ISPs share a single ASN. We first explain how we mapped ASNs to ISPs, followed by an explanation of how we collected data on IP addresses that were observed to facilitate spoofed traffic. Finally, we explain our methodology to obtain and analyze the datasets for network complexity, security effort, and institutional environment.

4.1 Mapping observations to ISPs

We define an ISP as a company that provides access services, typically in residential broadband markets. To map to ISPs our indicators and our observations of compliance, we need to identify their network address space.

We start our identification of the network space of ISPs with market analysis data from Telegeography: the GlobalComms database [45]. The database contains a highly reliable overview of the main broadband ISPs in each country, drawn from annual reports and market filings. We focused on the ISPs in 64 countries who together possess a broadband market share of over 85% in those countries [46]. This gives us a total population of 334 ISPs.

In the next step, we used CAIDA's AS ranking dataset [20]. It provides an approximate map of the organization name based on AS, the number of IP addresses announced, the country from which AS originate and the AS number. We then manually mapped the ASNs that belong to these ISPs by matching their names and the registration information to ASes that reside (at least partially) in that country.

In some cases, due to mergers, acquisitions, or branding changes, the AS name information might be outdated and no longer consistent with the current ISP name. The TeleGeography data also contains historical information about the ISPs. We search for historical names and updated mappings if we find evidence that an AS belongs to one of the ISPs in our dataset.

Finally, we look at the description of prefix announcements from the Hurricane Electric dataset [3] and exclude ASes that appear to be used primarily for other purposes like hosting, cellular data, IPTV, etc. There is a possibility that an ISP might provide multiple services from the same AS. In such a scenario, the identified AS might include some services like hosting infrastructure inside an access network. For our purposes, however, it still falls within the category of providing access services and should be included in the mapping of ISP network space.

We then map the IP addresses belonging to each AS number using BGP data from the Routeviews project [12]. Via the AS, we can then connect the IP address to the ISPs and country. Now, as some ASes span multiple countries, we geo-locate all IP addresses using the MaxMind GeoIP2 database [34]. For each ISP, we map only the portion of the AS that geo-locates to the country in which the ISP resides. This way, multi-country ASes get split up over the subsidiaries of the ISP in the various countries.

4.2 Data on IP Spoofing

To measure whether networks allow outbound spoofing to their upstream networks, we analyze data from the Spoofer project and from our Internet-wide scans of misbehaving open DNS resolvers. We merge these two sources into a variable that indicates non-compliance at the /24 prefix level. In this section, we first give an overview of the two techniques, followed by why and how of the data aggregation methodology.

Spoofer Project: The Spoofer project is the most known and used source to collect data on BCP38 compliance. The Spoofer tool is a client-server application. The client application is run by volunteers. It generates packets with spoofed and non-spoofed source addresses and then sends them to the Spoofer project server periodically and when it detects a new network. Based on the reception of these packets, the server infers whether the network blocks spoofed traffic or not. The benefit of these measurements is that it not only reveals networks that allow spoofed traffic to upstream networks, but it can also detect the opposite: networks that are compliant. However, data collection is based on volunteers to run the application from within the network, which limits the visibility of the tool across all ISP networks. This introduces some selection bias, where ISPs with more users as such and especially more users in the western countries, where the Spoofer project is more known, have higher odds of being included in the measurements. In this paper, we used data from the Spoofer tool collected over a period of 6 months (April-September 2019). The dataset contains tests collected from within 66 ISP networks in 31 countries. It is 26% of the total ISP population we have in our dataset.

Misbehaving Open DNS Resolvers: Jared Mauch first mentioned the idea to detect non-compliant networks using misconfigured open DNS resolvers on

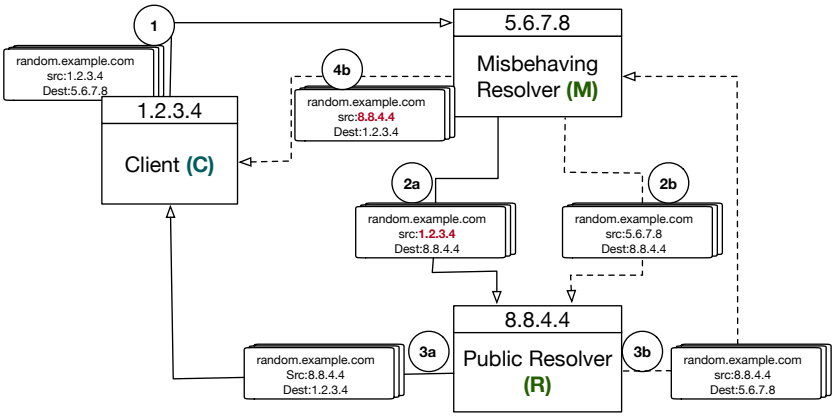

Fig. 2. Detecting spoofing ASes from misconfigured openresolvers.

the NANOG mailing list [33]. Subsequently, Kührer et al. scanned the IPv4 address space for misconfigured DNS resolvers [28] and found 2,692 ASes that allow spoofing.

It is important to note here that we are only interested in a very specific subset of open resolvers, namely CPE (customer-premise equipment) devices with a specific configuration error, which can function as a vantage point to observe the absense of SAV in parts of the network. In other words, these specific devices provide a de facto measurement platform for networks that lack SAV, since these devices respond with spoofed traffic in response to a specially crafted DNS request. In that sense, these misconfigured devices are similar to the spoofer client.

A previous study [28] has fingerprinted these misconfigured open resolvers and found the majority of them were running on home routers. These open resolvers have a very specific configuration error that makes them act as a forwarder for incoming packets. They probably have either misconfigured NAT rules or erroneous DNS proxy implementations [28]. We use these devices basically as vantage points. If we receive a response, this tells us two things: first, that there are no edge controls in place and, second, that there is no SAV at the network level, i.e., on border routers. In other words, there is no compliance with BCP38. In short, the misconfigured resolvers are a means for us to measure that anti-spoofing measures are not in place within the ISP's network.

Figure 2 illustrates how the proposed misconfiguration works. A client (C) with an IP `1.2.3.4` sends a DNS request to an IP address `5.6.7.8` to resolve `random.example.com` in step 1. If it reaches a misbehaving DNS resolver (M), it may forward the DNS request to another DNS resolver (R) to resolve the request in step 2a (in the example it is Google's open public resolver with an IP address `8.8.4.4`). However, a misbehaving resolver M may not change the source IP address of the original request to its own before forwarding it to R (step 2a). If SAV is not deployed at the network edge of the tested network, the forwarded query will reach R, which will perform the resolution and will send

the response directly to the client in step 3a, revealing the lack of SAV in the tested network. In a second scenario, the misbehaving resolver M may correctly forward the query to R by replacing the client's IP with its own in step 2b. After performing a recursive query resolution, R may send the response back to M, in step 3b. However, M may forward the DNS response to the client C without changing the source IP address of R to its own in step 4b, thus again revealing a non-compliant network.

To collect the data on networks without SAV using the method explained above, we extended the implementation of the myDig software [42]. We generated a list of unique subdomains for each routable IPv4 address and sent DNS requests from our server. Each time we get the response, we compare the destination IP address we sent the request to with the IP address that replied. We conclude that spoofing is feasible if the IP address from the response belongs to a different AS than we initially queried. We have repeated the scan for a period of 6 months from April till September 2019 on a weekly basis. With this technique, we observed IP addresses that can send spoofed packets in 240 ISPs and 60 countries.

From Measurements to Networks: The Spoofer and Open Resolver measurements observe individual IP addresses where sending spoofing traffic was feasible. A key methodological issue is how to infer, from these individual IP addresses, what overall portion of an ISP's network is non-compliant. In order to estimate the amount of IP space that allows spoofing, we need to aggregate the tested IP to the prefix level. SAV compliance requires configurations of the routers; hence it is more likely that either the entire prefix is compliant or not. It is, however, challenging to infer how the ISP has segmented its network in different prefixes, since it is operator dependent and is not reported publicly.

In principle, we could aggregate the non-compliant space at three different levels. First, we could classify the entire AS as spoofable if we find measurements showing that IPs can spoof in either of our two datasets. However, if policies are implemented on a prefix level, it will mean we would overcount the amount of addresses space that is non-compliant. Moreover, some large ASes operate across different countries, containing multiple ISPs (country-level subsidiaries of a multi-national ISP). Parts of the same AS might therefore be under the control of different organizations, which might result in varying SAV policies within that AS.

The second and more realistic approach would be to deduce policies from BGP inter-domain routing tables. The BGP table contains reachability information, which is shared amongst the ASes. We can map the IPs observed in our two datasets to the longest matching prefix from the BGP routing table and count that prefix as being non-compliant. These counts are better than assigning the entire space of ASN based on a few measurements. However, it still suffers from overcounting due to IP space aggregation by ISPs for efficient routing.

In this paper, we have chosen a more conservative third approach. If we find one or more IPs that allow spoofing within a /24 prefix (256 addresses), we

classify the entire /24 prefix as non-compliant. We also check BGP routing data and if an AS is announcing a prefix that is less than /24, we chose the smaller prefix as the more conservative estimate.

In [32], Luckie et al. show that in about 30% of the remediated cases in Spoofer, the client can still spoof address space outside the /24 prefix. This confirms that our approach is conservative and likely to underestimate the portion of an ISP's IP space that allows spoofed traffic to leave the network.

In Figure 3, we show the distribution of non-compliant IP space per ISP. It can be seen that 40% of ISP have a non-compliant address space of less than 1000 IP addresses, while around 86% percent has 10,000 or fewer IPs that can potentially send spoofed packets. In terms of /24 prefixes, we observe that around 16% ISPs have only one prefix, while around 37% of ISPs we have measurements from two /24 prefixes.

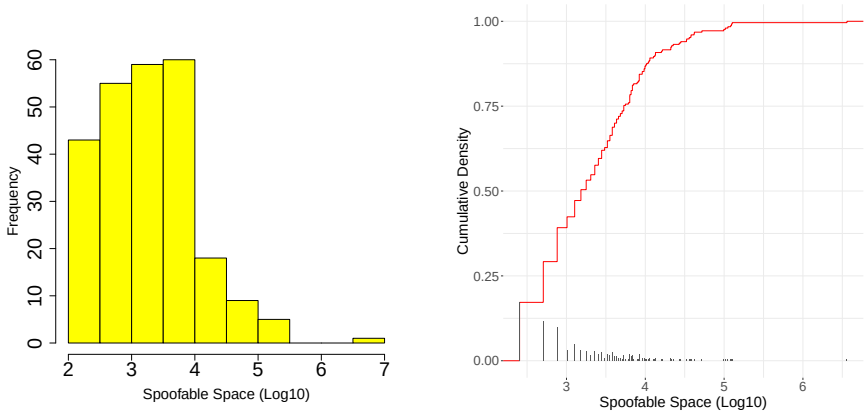

Fig. 3. Distribution of Amount of Non-Compliant Address Space of ISPs

Comparison of Datasets: The Spoofer client not only detects when a network allows a host to send spoofed traffic, but when it is blocked. The latter can be because the network has implemented BCP38 or because the client IP addresses are behind a NAT. We excluded the observations that detect a NAT, as it is unclear whether BCP38 is implemented at the network level or not. Prior work already found some networks with NAT not to be compliant when using IPv6 measurements [32]. We categorized IPs as "Spoofer Blocked", for which Spoofer infers the presence of BCP38.

In Figure 4, we summarize the overlap between Spoofer and Open Resolver datasets at the /24 prefix level – in other words, the number of /24 prefixes where we have observations from both datasets. There are only two /24 prefixes where the Open Resolver data has an observation and the Spoofer dataset contains an

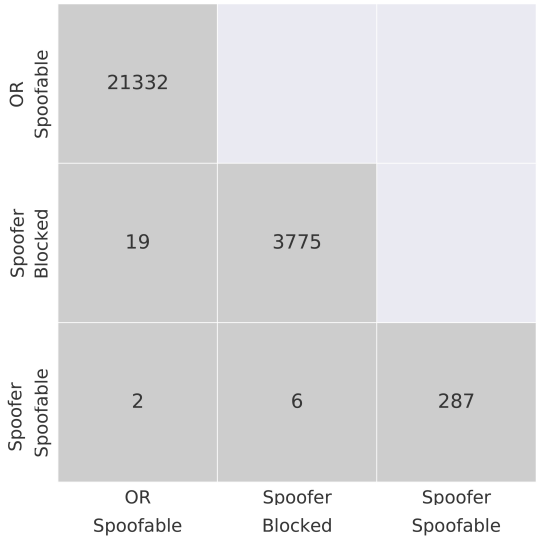

Fig. 4. Overlap among datasets at /24 level. Each number represents the number /24 prefixes where two datasets both have at least one observations. As a point of reference, we also include the overlap of each dataset with itself – i.e., the total number of prefixes for which that dataset has observations

observation that an IP address in that prefix allows spoofing (Spoofer Spoofable). We can attribute this tiny overlap to the lower coverage of the Spoofer tool compared to the Open Resolver dataset. Moreover, within the same /24 prefixes, we also find mixed results: 6 prefixes where Spoofer observes both spoofable and unspoofable traffic and 19 prefixes where Open Resolver observes spoofing is possible, but Spoofer finds it is blocked. Compared to the total dataset, this fraction of inconsistent results is negligible. They might result from differences in timing, where the prefix allowed spoofing during one observation and not during a subsequent one. Or they might result from the fact that SAV is implemented at a smaller prefix level than /24. As we discussed above, [32] found that in 90% of the cases, operators implement SAV at the /24 level or larger. That still leaves 10% where operators implement it at smaller prefixes, which might result in different SAV policies in the same /24. We do not observe any contradictory test results for the same IP address.

Again, these inconsistencies occur in a tiny fraction of our overall observations. Even in these cases, we have proof that spoofing is allowed from a least a portion of the /24. For this reason, we consider this fraction of prefixes also to be non-compliant. In the rest of the paper, we only consider the non-compliant address space.

Figure 5 shows where we have observations for the same ISP in two datasets. ISPs typically operate their address space as multiple networks, sometimes even multiple ASes. This explains why a large number of ISPs where we detected

spoofing also had other parts of their network where SAV was implemented. It is important to note that, on average, we find 10.5 more non-compliant addresses than compliant addresses in these ISPs. In any case, as explained earlier, even if there are a few addresses that allow spoofing, the possibility of a DDoS attack remains intact. From our overall population of 343 ISPs, we have observations that indicate non-compliance for 250 of them (73%). For 182 ISPs (53%), we also have test results indicating that they did deploy SAV on some prefixes. We have no test results whatsoever for 51 (15%) of ISPs. For 149 of all ISPs (43%), we have mixed results: hosts can send spoofed packets from some prefixes and while spoofed traffic is filtered in other prefixes.

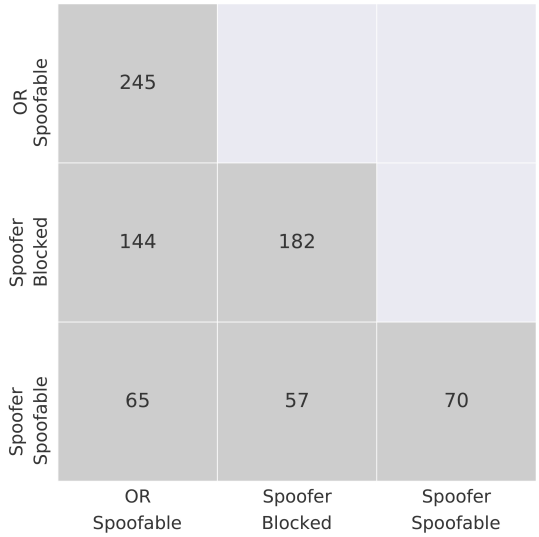

Fig. 5. Overlap among datasets at ISP level. Each number represents the number of ISPs both datasets have at least one observation. As a point of reference, we also include the overlap of each dataset with itself – i.e., the total number of ISPs for which that dataset has observations

4.3 Network Properties Data

Total Size of Advertised IP Space: We used routing data collected by Routeviews project [12] to estimate the number of IPs per ISP. We analyzed weekly BGP routing data for the period of April-September 2019. We used the pyasn library [14] to determine the prefixes announced for the ASNs, which are then mapped to ISPs in our dataset. We then aggregated the total number of IPs for each of the ASNs per ISP. Finally, we took the average number of IPs from the weekly data per ISP.

Average Number of Prefixes: IP space is announced in BGP tables in the form of prefixes. An IP prefix represents the number of bits, which is used to identify a network and determine the total number of hosts. Network operators usually aggregate the total advertised space to the maximum prefix announcement possible for efficient routing and lower number of advertisements, which is useful routing tables. However, in some cases, due to routing policies or their usage, they advertise multiple prefixes for the same range. Once a week, we calculated the **number of prefixes** announced per ISP between Apr-September 2019. We then took an average of these counts.

Stability of Prefix Announcements: Moreover, using the BGP data and pyasn, we determined the **stability** of the announced prefixes by calculating the percentage of prefixes per operator that remained unchanged compared to the total set of prefixes that were announced at some point during the measurement period (Apr-Sep, 2019).

4.4 Security Effort Data

RPKI: Internet operators use BGP to exchange routing data. It contains prefixes and the number of hops they are away from the AS announcing the information. Routing information is constantly changing and BGP is flexible enough to converge for these route changes. However, BGP lacks a mechanism to validate if the prefixes being announced actually belong to the entity announcing them. To verify the authenticity of the announcement, Internet Engineering Task Force (IETF) developed a mechanism known as Resource Public Key Infrastructure (RPKI) [19]. Internet operators can now use a cryptographic system of public/private keys to sign prefixes, thereby authenticating that they are authorized to announce them. The Regional Internet Registries (RIRs) maintain public key certificates. The operators can detect announcements with an invalid route origin. It is important to note that RPKI doesn't secure the path. It is, however, the first step towards BGP route security.

We interpret the adoption of RPKI as an indicator of security effort by the ISP. We used Nlnetlabs Routinator, an RPKI validator tool [10], to download prefixes that were signed by their respective ISPs. We then mapped these ASNs, where we observed signed prefixes to ISPs from our dataset. We assigned a binary value of 1 to the ISPs that had signed one or more prefixes, and 0 to the ISPs that had no signed prefixes.

Spam Bots: Like the lack of SAV compliance, infected end-user machines in an ISP network are a widely recognized security externality [16]. While ISPs have been involved in mitigating botnets, many of the benefits of mitigation go to the third parties that are attacked by botnets. Contrary to SAV, though, the ISP might suffer some cost via blacklisting when it does not mitigate outbound spam. We interpret the relative number of bots is an indicator of the security effort ISPs are willing to undertake in light of significant external effects.

We measured the number of bots in the ISP networks using multiple data feeds. First, we used the Composite Blocking List (CBL) of SpamHaus [43]. The dataset contains IPs of spam bots, including Cutwail, Rustock, Lethic, Kelihos, and Necurs. We receive a daily report from SpamHaus and map the IPs in the feed to their respective ISPs. We have to control for DHCP churn, which would lead to serious overcounting for ISPs with very dynamic IP address allocation, as the same infected machine would show up under multiple IP addresses. [38]. In order to compensate for this churn, we count the number of unique IPs observed each day and then calculate the average of all daily counts over our measurement period. The downside of this methodology could be that we are undercounting. However, due to our long time-frame, the averages would bring the estimated count closer to the actual number of infected machines.

A second indicator on bots is based on a spam trap operated by Dave Rand of TrendMicro. We follow the same approach as with Spamhaus CBL: extract the IP addresses, map them to the ISPs and calculate the daily average number of unique IPs seen over the measurement period.

Amplifiers: Finally, we look at the presence of so-called amplifiers in the networks of ISPs. Amplifiers are legitimate services that can be abused in amplification attacks with spoofed traffic, exactly as was explained at the start of this paper. Again, this is an example of a security externality for the ISP, thus providing us with an indicator for measuring security effort related to threats with significant external effects.

We downloaded Rapid7 data containing IP addresses of UDP amplifiers in ISPs' network [39]. Rapid7 scans for various protocols are publicly available every month. In our paper, we used IP addresses for Chargen, DNS resolvers, Memcached, Netbios, Ntpmonlist, Portmap and Qotd.

We have decided to combine the observations of amplifiers into a single proxy. Our goal is to capture a signal on overall network hygiene, not the ISP response to a specific type of amplifier. The protocols that we have included in our study have been identified as potential attack vectors by US CERT advisory [11]. Network operators should either take down the amplifiers or at least deny access to the services over the Internet. They can also deploy Response Rate Limiting (RRL) to reduce the rate at which replies are sent and thus limiting the impact of amplification. By combining the amplifier observations, we also get much better coverage of observations across the ISP population, further improving the statistical behavior of the proxy. (To telegraph ahead to the statistical analysis: when we include each amplifier type as a separate predictor in the model, the sigal gets too weak and we no longer find any significant relationships.) A high correlation between non-compliant networks and the presence of a large number of amplifiers would indicate operators' inaction for DDoS problem. We mapped the reported IP addresses to each ISP. Finally, to mitigate the effect of churn, we calculated daily averages of the number of observed amplifiers using the methodology explained above.

4.5 ISP Characteristics Data

Number of Subscribers: We used the total number of subscribers as a proxy to determine the size of ISPs. Telegeography database reports the total number of subscribers per quarter, and we selected quarter two of 2019, as it matches the closest to the spoofing datasets.

Average Subscription Price and Revenue: Telegeography reports revenue and average subscription price per company. However, they do not have data for all Internet providers. We mapped revenue and subscription prices to our dataset. We are missing 68 ISPs and do not have any reliable estimates to fill in missing values.

4.6 Institutional Environment Data

Finally, we collected indicators for the institutional environment of the ISPs. The first one is at the level of countries, the second at the level of the provider community, i.e., whether the ISP is part of a group of industry peer committing themselves to adopt good security practices for routing, among which is SAV.

ICT Development Index: We also used the ICT development index, which is an indicator representing ICT development per country. The dataset is provided by ITU (United Nations International Telecommunication Union). It assigns values from 1 to 10, with a higher number representing a higher level of development based on various ICT indicators.

MANRS Dataset: MANRS initiative requires best practices for ISPs to reduce the most common routing threats. We downloaded member ASNs of MANRS from their website and mapped it to ISPs in our dataset [5].

5 Statistical Model for Non-Compliance

In this section, we first explain the transformations we did for some of the indicators, followed by basic statistics of the dataset. Next, we estimate a linear model and discuss the results and interpretations.

The number of non-compliant IP addresses per ISP has a correlation of 0.52 with the total number of IP addresses being announced by that ISP. ISPs with a larger number of IP addresses have a higher chance of having tests and are hence more likely to have non-compliant address space being observed. For this reason, we first transform the dependent variable into a relative metric: the ratio of the non-compliant address space to the whole address space being announced by the ISP. In Figure 6 (a), we show the distribution of this normalized variable. One of the concerns is that the distribution is left-skewed, partially because of the fact that we adopted a conservative approach to estimate the amount of

non-compliance address space, likely undercounting it. This distribution would violate the assumptions of linear regression. We perform a natural log transformation to resolve this issue. Figure 6(b) shows that the transformed distribution is much closer to normal. We used this transformed variable as the dependent variable in our model. For the same reasons, we also log-transform some of our independent indicators, namely the number of subscribers and the security effort indicators for bots and amplifiers. Table 1 summarizes the indicators that are used in the model.

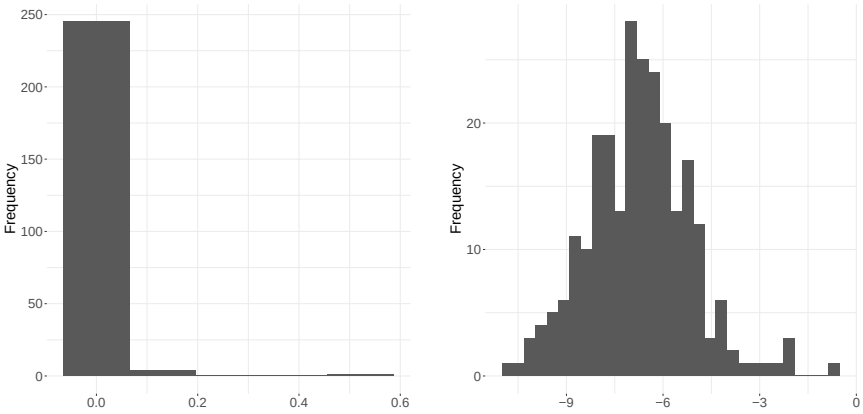

(a) Portion IP Space Non-Compliant (b) Log-transformed Portion IP Space Non-Compliant

Fig. 6. Distribution of Transformed Spoofable IPs per ISP

5.1 Model Specification

To reiterate: we measure SAV compliance as the number of IP addresses in all /24 blocks where one or more IP addresses were observed as being non-compliant in our Spoofer or Open Resolver datasets. Our response variable is a normalized count of non-compliant addresses divided by the total number of IP addresses announced by the operator. We define our response variable Y_i as the log of the normalized size of non-compliant address space by ISP_i for $i = 1, \ldots, n$, where n is the total number of ISPs for which we have our tests. To estimate the impact of network properties, security effort, ISP characteristics and institutional factors on non-compliant address space, we use a linear regression, which takes the following form:

$$\ln(Y_i) = \beta_0 + \sum_{j=1}^{k} \beta_j x_{ij} + x_{ij} * y_{ij} + \epsilon_i$$

where β_0 is the intercept and x_{ij}, $j = 1, \ldots, k$, are the indicators for network complexity, security efforts, ISP characteristics and institutional environment and $x_{ij} * y_{ij}$ is the interaction terms for the model. Our error term ϵ_i is normally distributed with mean 0 and the variance sigma squared.

Variables	**N**	**Min**	**Mean**	**Median**	**Max**
Portion of adv. IP space non-compliant	250	0.00002	0.006	0.0012	0.522
Total size of adv. IP space	250	6,617	6M	1,4M	124M
Avg # of prefixes	250	3.61	204.0	88.15	2,668
AS stability (percentage)	250	4.03.	79.75	86.38	100
Avg # of bots Spamhaus	250	18.3	29,772	4,264	1,1M
Avg # of amplifiers	250	86.3	13,979	2,703	0,51M
Avg # of bots spam trap	250	1.06	43.25	14.36	1,681
RPKI	250	0	n/a	1	1
# Subscribers	250	5,500	3M	0,7M	174M
Avg sub price (USD)	182	4.5	47.02	41.04	883
Revenue (USD)	182	12	2,302	533.7	46,377
ICT Dev Index	250	2.42	6.68	7.04	8.98
MANRS	250	0	n/a	0	1

Table 1. Summary of indicators used in the model

5.2 Discussion of Results

We construct models following step-wise inclusions for the various indicators. A summary of four models is presented in Table 2. Our goal is to understand the relationship of various indicators and improve the goodness of fit for these models. The Adjusted-R-squared value increases from 0 in the model (1) to 0.47 in the model (3), which means we were able to explain 47% of variance by adding the indicators for network complexity, security effort, and institutional environment. Moreover, the signs of coefficients do not change from the model (2) to model (4). Note that we had to drop 68 ISPs in the model (4) due to missing information on revenue and subscription prices. Since these two indicators do not add any explanatory power, nor affect the other variables, we omit them and select model (3) as our final model. It performs the best in terms of explained variance.

In the final model, we also include interaction terms to understand the effect of advertised IP space in combination with subscribers and prefixes. Note that the distribution of IPv4 address space is asymmetric: early adopters were able to acquire a large number of addresses, while more recent market entrants got smaller allocations for the same number of users – due to the shrinking pool of address space held by RIRs – and they therefore have to rely more on NAT. Furthermore, IPs are allocated to ISPs in terms of prefixes. The early adopters were able to acquire bigger ranges. Later on, ISPs were allocated smaller prefixes,

	Response Variable: Portion of adv. IP space non-compliant (ln)			
	(1)	(2)	(3)	(4)
# Advertised IPs (ln)		-1.782^{***}	-2.080^{***}	-1.844^{**}
		(0.342)	(0.382)	(0.617)
Avg # prefixes			0.011^{**}	0.016^{***}
			(0.003)	(0.004)
AS stability (percentage)			0.001	0.002
			(0.004)	(0.004)
RPKI			-0.313^{*}	-0.207
			(0.156)	(0.187)
Avg # amplifiers (ln)			0.368^{***}	0.283^{***}
			(0.075)	(0.080)
Avg # bots SH.(ln)			0.099	0.071
			(0.077)	(0.090)
Avg # bots ST.(ln)			-0.068	0.027
			(0.090)	(0.099)
# Subscribers (ln)		-0.842^{*}	-1.257^{**}	-0.936
		(0.353)	(0.410)	(0.642)
Average subs price				0.001
				(0.001)
Revenue				0.00001
				(0.00003)
ICT Dev Index			-0.164^{*}	-0.236^{**}
			(0.065)	(0.081)
Adv. IPs(ln):# subs(ln)		0.077^{**}	0.094^{***}	0.079
		(0.024)	(0.028)	(0.045)
Adv.IPs(ln):Avg# prefixes			-0.001^{**}	-0.001^{***}
			(0.0002)	(0.0003)
Constant	-6.725^{***}	15.000^{**}	18.924^{***}	14.959
	(0.101)	(4.736)	(5.205)	(8.555)
Observations	250	250	250	182
R^2	0.000	0.333	0.495	0.463
Adjusted R^2	0.000	0.325	0.472	0.421
Residual Std. Error	1.601 (df = 249)	1.315 (df = 246)	1.163 (df = 238)	1.134 (df = 168)
F Statistic		40.929^{***} (df = 3; 246)	21.209^{***} (df = 11; 238)	11.131^{***} (df = 13; 168)

Note: *p<0.05; **p<0.01; ***p<0.001

Table 2. Linear regression model

as IPv4 started running out. This presents an interesting interaction effect: ISPs with a large number of IP addresses and a small number of prefixes would have a relative advantage in SAV adoption because routes would be easier to configure and maintain. We included the interaction of total IP space advertised with the number of prefixes to test this hypothesis.

The indicators that measure the size of the ISP, in terms of address space and in terms of subscribers, are significant and have a negative sign. However, we need to be careful with their interpretations due to interaction effects. We explain this in more detail later in the section, but to telegraph ahead: we observe that the signs change for both the number of advertised IPs and the number of subscribers when the number of announced IPs get past the 450,000 mark (see intersecting point Figure 6). From our dataset, a large majority (76%) of ISPS advertise more than 450K IP addresses. We can therefore state that, by and large, larger ISPs have a higher proportion of non-compliant IP space.

In the case of BCP38, size plays an essential role in the implementation of SAV. ISPs with larger address space are more likely to peer with a higher number of upstream providers, to avoid a single point of failure. To be compliant, they would then have to implement BCP38 on multiple edge routers. This would be more costly. However, a counteracting effect of size, especially when measured in terms of the number of customers, is that larger ISPs have more resources and expertise than smaller operators. Furthermore, there are likely economies of scale in implementing BCP38. The model suggests these cost-reducing effects of size are smaller than the cost increases because of the increased network complexity. In the case of the number of prefixes, we observe for around 96% of ISPs increasing the number of prefixes would also lead to an increase in non-compliant IP space. We give a detail explanation in the latter part of the section. We did not find the stability of prefixes significant in our model.

Next, we look at the impact of ISP security efforts. We have used the signing of BGP prefixes (RPKI), as a positive indicator of effort, i.e., the willingness to invest in security issues with significant externalities. The number of DDoS amplifiers and spambots in the network is a negative indicator of this willingness. The model finds a weak but significant relationship with RPKI. Operators that sign their prefixes are more likely to implement SAV in their networks. From our results, holding all variables constant, an ISP that signs its prefixes with RPKI will have a 31.3% lower portion of non-compliant space compared to ISPs that don't sign their prefixes. We find the indicator for the number of amplifiers per ISP has a significant positive impact. In other words, for a 1% increase in the number of amplifiers, there is an increase of 0.36% of the portion of non-compliant address space, holding all other variables constant. Please note that when we treat each amplification protocol independently in the model, the relationships are no longer significant. However, by combining these observations, we can approximate overall hygiene, as observed by the fact of whether operators reduce amplifiers across the board.

The two indicators for spambots are not significant. There could be several reasons for this. Contrary to RPKI and the number of amplifiers, the number of

spambots is influenced by attacker behavior. This could confound the indicator in terms of measuring provider effort. Also, ISPs have another incentive than security for dealing with spambots: they might get blacklisted. This, in turn, might impact the service quality for their customers (e.g., legitimate email might also get blocked). In other words, this indicator might also include effects that are not capturing the provider's willingness to invest in security issues with serious externalities.

When looking at the ISP characteristics, we find that the number of subscribers has an impact. We discuss this below, where we interpret the interaction effects. The other two indicators, average subscription price and revenue of the ISP, were included in the model (4). Both of these variables are non-significant and have a small coefficient. This might be partially due to missing data, since we have no information for about 27% (68) of ISPs. Future work is needed to collect data on the financials of ISPs to understand these relationships.

Next, we measure the impact of the institutional environment. The model shows a weak but significant effect for the ICT development index. ISPs that operate in countries with lower ICT development have a higher percentage of non-compliant address space. In other words, for a 1% increase in ICT index, there is a decrease of 16.4% of the portion of non-compliant address space, holding all other variables constant. We did not regress our model against the MANRS indicator since we only found 16 ISPs out of 250 MANRS signatories in our data on non-compliance. This makes sense, as we are only looking at ISPs that have been observed as allowing spoofed traffic to leave their networks. MANRS signatories explicitly commit to adopting BCP38. It seems that this self-commitment does have an effect, but we would need more test results from compliant networks to confirm this. Here, all we can say is that most ISPs that signed MANRS are not observed as allowing spoofing.

We interpret the impact of the number of subscribers and prefixes in more detail. As these variables are influenced by the amount of announced IP space, we included 2 interaction terms in the model. The coefficient of the interactive term (Adv. IPs(ln):# subs(ln)) is positive and statistically significant at .001 level. On the other hand, the interactive term (Adv. IPs(ln):Avg # prefixes)) is negative but also statistically significant at .01 level. This tells us that the coefficient of announced IPs depends on the value of subscribers and prefixes and vice versa; these estimated coefficients are conditional. It does not indicate anything, however, about the magnitude or statistical significance of these conditional coefficients. To help understand the effects of this marginal coefficient, we plot in Figure 7 the relationship between the variables involved in both interaction terms.

The left plot in Figure 7 clearly shows that with increasing the announced IP space, the magnitude of the coefficient of the number of subscribers also increases, ranging from a -0.42 for the minimum number of announced IPs to 0.50 for the maximum number of announced IPs. This means that when the number of IPs is lower than 450k (see the intersecting point), an increase in the number of subscribers will lead to a decrease in the portion of non-compliant IP

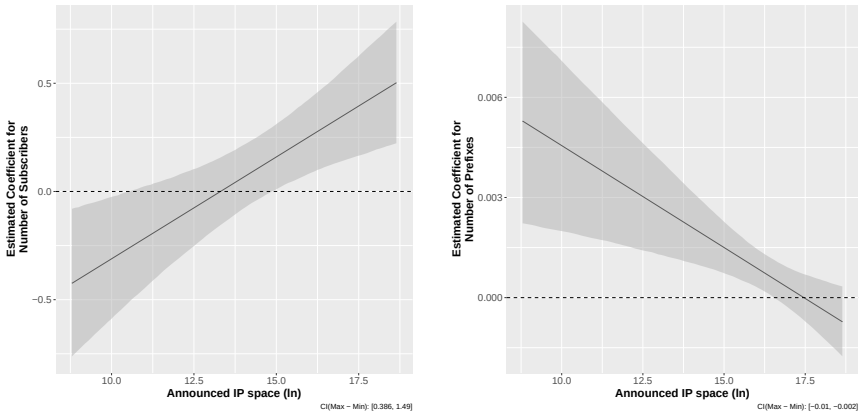

Fig. 7. Interaction of Subscribers and Prefixes with the total number of IP addresses Announced by ISPs

space. On the contrary, for those ISPs with more than 450k advertised IPs, an increase in the number of subscribers would lead to an increase of the portion of non-compliant IP space. For instance, for the ISP with the largest number of advertised IP addressed, a 1% increase in the number of subscribers will increase the portion of spoofable IPs by 0.5%. The confidence intervals (see caption at the right bottom corner of Figure 7) of the difference between the conditioned effects of the announced IP space at the minimum and maximum values of the subscribers.

The second interaction term shows a negative relationship between conditional coefficient of the announced IP space and the number of prefixes, i.e., increasing the announced IP space, the magnitude of the coefficient of the number of prefixes decreases, ranging from a 0.05 for the minimum number of announced IPs to -0.0007 for the maximum number of announced IPs. This means that when the number of IPs is larger than 39 million, an increase in the number of prefixes will lead to a decrease of the portion of non-compliant IP space. On the contrary, for those ISPs with less than 39 million advertised IPs, an increase in the number of prefixes will lead to an increase of the portion of non-compliant IP space. For instance, for the ISP with the largest number of advertised IP addressed, a 1% increase in the number of prefixes will increase the portion of spoofable IPs by 5%. From our dataset, we have 242(96%) of ISPs that advertise less than 39 million IPs.

In summary, we observe that network complexity plays a significant role in non-compliant IP space. Our interaction terms for the number of IPs announced with subscribers, and prefixes give us insights on the variability of policies for significantly large ISPs. From our security effort indicators, we found the number of amplifiers significant, which shows us that non-compliant ISPs are also part of the bigger ecosystem of DDoS attacks since amplifiers are commonly

used to redirect and amplify the spoofed IP packets. Other than the number of subscribers, we did not find the rest of the ISP characteristics significant.

Comparing this with institutional factors, ISPs that are in countries with better ICT infrastructure are more compliant. We observe very few ISPs within MANRS, which is a positive sign. However, future work is needed to understand the role of MANRS for compliance.

6 Challenges in the Adoption of SAV

Lack of SAV by network operators has been a concern for many years now. In a recent survey by RIPE NCC, the RIR for Europe, West Asia and the former USSR, 4,161 operators responded that DDoS was the most significant security problem for them [9]. Even though ISPs acknowledge that DDoS is a considerable challenge, we still find evidence of non-compliance in part of the networks of 250 ISPs (73%). A key reason is the cost associated with the adoption and maintenance of SAV, while at the same time providing very limited benefits for ISPs. Moving forward, we need to re-align the incentive structure if we want to see any uptick in compliance. In this section, we review some of the available options and the role various actors can play to improve SAV compliance and reduce the number of hosts that can successfully send spoofed traffic and launch amplification DDoS attacks.

6.1 Reducing the cost of adoption

One of the ways to persuade operators to comply is by reducing the cost of SAV adoption. Our empirical findings emphasize the importance of cost incentives for smaller providers. They have, on average, higher rates of non-compliant address space. Larger providers benefit from the resources, non-compliance in their network can largely be attributed to limited incentives. However, it is challenging for smaller ISPs to configure SAV correctly for their network.

In the RIPE survey, 36% of the respondents suggested the best way RIPE can help them is by providing security-specific training programs. It is important to note here that RIPE attracts participants from various businesses, including ISPs, hosting providers, educational networks, etc. The security issues they require help with do not necessarily all apply to ISPs. In another survey among 84 network operators, a vast majority of respondents reported that SAV is out of reach for them in terms of knowledge, planning, and time need to maintain up-to-date access control lists (ACLs) for the implementation of SAV [29].

There are multiple ways network operators can get support to implement SAV, including training engineers via industry communities like RIPE and Network Operators' Groups (NOGs). Another option is getting router vendors to providing SAV compliance as the default option [32]. Other areas in security have shown that better tools can also reduce the cost of adoption of new solutions. Notable examples include an RPKI prefix signing tool by RIPE. It provides a

simple web-based or API based interface for providers to sign prefixes [8]. Similarly, Let's Encrypt, a non-profit certificate authority, has played a significant role in improving HTTPS adoption after major browsers started flagging HTTP sites as insecure. It is currently serving over 180 million websites [4]. One of the key reasons for this success is open-source, free software with clear and concise documentation that requires few clicks or commands to configure a TLS certificate to serve HTTPS traffic.

A usable open-source tool for SAV implementation would have to accommodate the dynamic nature of Internet routing. Building such software is a challenging task, due to complex routing policies based on the various needs and contractual relationships of ISPs. Moreover, the tool would be required to keep updated information about customers and network allocations to feed automated systems. Any mishap potentially disruptive and cause downtime for customers. Unless thoroughly tested and backed by major players, it is highly unlikely that ISPs would use a tool that affects the backbone of their business.

Some other suggestions include decentralizing the BGP routing [13] or offloading it to cloud where SAV implementation could benefit from economies of scale [7]. Recently, the University of Massachusetts Amherst has received a $1.2 million grant to develop and test "logically centralized interdomain routing architecture." [7]. It is yet to be seen how this solution will pan out and if Internet providers would trust cloud networks to route their traffic. Another new direction is the emergence of reconfigurable networks [18]. The P4 programming language, in combination with supported hardware, would enable network operators to change the configuration of the connected switches without any downtime. It is still in its early days, and it is hard to predict if operators will end up adopting it.

All in all, some options to reduce cost are within reach, such as training via industry associations, but substantial cost reductions would require some form of automation. This still seems infeasible in the short term.

6.2 Reducing Information Asymmetry

Another way to re-align the incentives is to reduce information asymmetry. In a recent development, Cloudflare launched a website isbgpsafeyet.com to reduce information asymmetry on RPKI deployment. Participants can run a test to check if their ISP would accept a legitimate route with an invalid announcement. If the test fails, it means the ISP would likely accept a leaked or a hijacked route. The users are encouraged to tweet the results to increase awareness, which might result in increasing pressure on ISPs to implement RPKI.

CAIDA's Spoofer project has been publishing lists of non-compliant and remediated prefixes and ASes on its website. It sends updates to various network operators mailing lists making transparent which networks are not compliant and which networks remediated. However, they do not present results at the level of ISPs, but at the level of ASes and prefixes, i.e., technical identifiers rather than actors. AS names might not even match ISP names in some cases. Furthermore,

these results do not rank networks in terms of compliance. If there is a reputation effect or social proof nudge to be gained from making non-compliance and remediation more visible, then it would likely be more effective at the level of the ISP, since that is the actor who needs to be incentivized to remediate. Similarly, ISPs that remediate non-compliant address space could be incentivized by receiving recognition in the industry.

Moreover, after the launch of the 'Is BGP safe yet' website, network operator groups (NOGs) (e.g., [6]) are already discussing to add a social nudge to www.bcp38.info. It is important to note that residential IP space provides an easily accessible vantage point for attackers to send spoof packets. If subscribers request their ISPs to be compliant, it will offer them an incentive to deploy SAV. We propose creating a more visible list of ISPs and their degrees of compliance to be shared not only on the operator mailing list but also with the users and on the website of the RIRs and perhaps with national CERTS.

6.3 Internalizing Externalities

A third option would be to internalize some of the external costs of non-compliance to the ISP. One study suggested regulating government procurement to require all government-contracted network providers to adopt SAV [32]. Another route might be pressure from other providers. They could, for example, require SAV to be implemented before entering into a peering agreement. Upstream providers, commonly known as tier 1 or tier 2 providers, hold a strong vantage point to observe SAV compliance, since they give connectivity to many ISPs. They can detect if ISPs are not filtering traffic, especially when they are the only provider of that ISP [31]. There have been examples where upstream providers have leveraged their position to achieve security improvements in BGP routing. Hurricane Electric and Portugal's IPTelecom joined forces to cut off Bitcanal from the global Internet, after it had consistently observed to conduct BGP route hijacking [1]. The organization was later also removed by German Internet exchange DE-CIX and others in the routing ecosystem.

6.4 Community Action to Reduce Weak Links

We see no path towards overcoming the weakest-link problem, i.e., the fact that a single non-compliant provider would mean amplification DDoS attacks are still possible. Miscreants who want to conduct such attacks need only to rent hosts in these non-compliant networks. From there, they can reach all amplifiers with their spoofed packets. That being said, the community of ISPs might reduce the number of such networks via collective action. MANRS is a good example of this. We could not include MANRS in the regression model precisely because operators that signed MANRS are unlikely to be observed in the tests of Spoofer and Open Resolver used in our study. In other words, participating in such an initiative does seem strongly correlated with SAV adoption. However, we did observe a small number of member ISPs (16 out of 250) with one or more prefix being non-compliant. While MANRS promotes best-practices and helps

with running a healthy BGP environment, operators are not legally bound to implement any of the promoted policies.

Moreover, some pressure can be exerted by National CERTs and RIRs on ISPs to behave more in line with community norms. We recently saw an example of the American Registry for Internet Number (ARIN) banning Cogent, a large ISP, from accessing its WHOIS database for 6 months, after several ISPs complained that they had received unsolicited marketing calls from the Cogent's sales team [2].

7 Conclusions

We have presented the first study to combine two independent techniques to measure the state of SAV across networks. We used these measurements to estimate the extent to which the population of 334 ISPs in 61 countries is not compliant with BCP38. A large portion of them, 73% to be precise, has at least one prefix that allows IP spoofing. What portion of its IP address space is not compliant is influenced by network complexity, security efforts, ISP characteristics and the institutional environment. As SAV adoption suffers from misaligned incentives, the main route forward seems to be to either reduce the cost of compliance for the providers or to increase the cost of non-compliance. Some of these forces seem to be at play, already, as a significant portion of the ISP population has in fact adopted SAV for all, or at least most, of its address space. Our study only looked at ISPs with at least one prefix that allows spoofing. We did not study the factors that explain why some ISPs are, in fact, fully compliant. Future work might look at that glass-half-full part of the picture and find factors apparently overrode the incentive misalignment that has plagues SAV adoption for a long time now.

Acknowledgments: The authors would like to thank Matthew Luckie, Radu Anghel, anonymous reviewers, and our shepherd Richard Clayton for their valuable feedback. This work was partially supported by the PrevDDoS project funded by the IDEX Université Grenoble Alpes "Initiative de Recherche Scientifique (IRS)" and the Grenoble Alpes Cybersecurity Institute CYBER@ALPS under contract ANR-15-IDEX-02.

References

1. BGP hijacker booted off the Internet's backbone, www.theregister.co.uk/2018/07/11/bgp_hijacker_booted_off_the_internets_backbone
2. Cogent cut off from ARIN Whois, www.theregister.co.uk/2020/01/09/arin_boots_cogent
3. Hurricane Electric Toolkit, https://bgp.he.net/
4. Let's Encrypt, https://letsencrypt.org/
5. MANRS, https://www.manrs.org/isps/participants/
6. NANOG, https://mailman.nanog.org/pipermail/nanog/2020-April/107356.html
7. NOIA Network, https://noia.network/technology

8. Resource Public Key Infrastructure (RPKI), https://www.ripe.net/manage-ips-and-asns/resource-management/certification/resource-certification-roa-management
9. RIPE NCC Survey 2019, https://ripe79.ripe.net/presentations/89-RIPE-NCC-Survey-2019-Report-Presentation.pdf
10. Routinator 3000, https://www.nlnetlabs.nl/projects/rpki/routinator/
11. UDP-Based Amplification Attacks, https://www.us-cert.gov/ncas/alerts/TA14-017A
12. University of Oregon Route Views Project, http://www.routeviews.org/routeviews/
13. Using Cloud Resources to Dramatically Improve Internet Routing, https://www.umass.edu/newsoffice/article/using-cloud-resources-dramatically-improve
14. Asghari, H.: pyasn–Python IP address to autonomous system number lookup module. URL: https://github. com/hadiasghari/pyasn
15. Baker, F., Savola, P.: Ingress Filtering for Multihomed Networks. RFC 3704 (Mar 2004), https://rfc-editor.org/rfc/rfc3704.txt
16. Bauer, J.M., Van Eeten, M.J.: Cybersecurity: Stakeholder incentives, externalities, and policy options. Telecommunications Policy **33**(10-11), 706–719 (2009)
17. Beverly, R., Bauer, S.: The Spoofer Project: Inferring the Extent of Source Address Filtering on the Internet. In: USENIX Steps to Reducing Unwanted Traffic on the Internet (SRUTI) Workshop (Jul 2005)
18. Bosshart, P., Daly, D., Gibb, G., Izzard, M., McKeown, N., Rexford, J., Schlesinger, C., Talayco, D., Vahdat, A., Varghese, G., et al.: P4: Programming protocol-independent packet processors. ACM SIGCOMM Computer Communication Review **44**(3), 87–95 (2014)
19. Bush, R.: Origin Validation Operation Based on the Resource Public Key Infrastructure (RPKI). IETF RFC7115 (January 2014) (2014)
20. CAIDA: AS Rank IPv4, https://asrank.caida.org/
21. CAIDA: The Spoofer Project, https://www.caida.org/projects/spoofer/
22. The Closed Resolver Project, https://closedresolver.com
23. ITU: http://www.itu.int/net4/ITU-D/idi/2017/index.html
24. Korczyński, M., Nosyk, Y., Lone, Q., Skwarek, M., Jonglez, B., Duda, A.: Don't Forget to Lock the Front Door! Inferring the Deployment of Source Address Validation of Inbound Traffic. In: Passive and Active Measurement Conference (PAM) (2020)
25. Korczyński, M., Nosyk, Y., Lone, Q., Skwarek, M., Jonglez, B., Duda, A.: Inferring the Deployment of Inbound Source Address Validation Using DNS Resolvers. In: The Applied Networking Research Workshop 2020 (ANRW) (2020)
26. Korczyński, M., Tajalizadehkhoob, S., Noroozian, A., Wullink, M., Hesselman, C., van Eeten, M.: Reputation Metrics Design to Improve Intermediary Incentives for Security of TLDs. In: 2017 IEEE European Symposium on Security and Privacy, EuroS&P. pp. 579–594 (2017)
27. Korczyński, M., Wullink, M., Tajalizadehkhoob, S., Moura, G.C.M., Noroozian, A., Bagley, D., Hesselman, C.: Cybercrime after the sunrise: A statistical analysis of DNS abuse in new gtlds. In: Proceedings of the 2018 on Asia Conference on Computer and Communications Security, (AsiaCCS 2018). pp. 609–623 (2018)
28. Kührer, M., Hupperich, T., Rossow, C., Holz, T.: Exit from Hell? Reducing the Impact of Amplification DDoS Attacks. In: USENIX Conference on Security Symposium (2014)

29. Lichtblau, F., Streibelt, F., Krüger, T., Richter, P., Feldmann, A.: Detection, Classification, and Analysis of Inter-domain Traffic with Spoofed Source IP Addresses. In: Internet Measurement Conference. ACM (2017)
30. Lone, Q., Luckie, M., Korczyński, M., Asghari, H., Javed, M., van Eeten, M.: Using Crowdsourcing Marketplaces for Network Measurements: The Case of Spoofer. In: Traffic Monitoring and Analysis Conference (2018)
31. Lone, Q., Luckie, M., Korczyński, M., van Eeten, M.: Using Loops Observed in Traceroute to Infer the Ability to Spoof. In: Passive and Active Measurement Conference (2017)
32. Luckie, M., Beverly, R., Koga, R., Keys, K., Kroll, J., claffy, k.: Network Hygiene, Incentives, and Regulation: Deployment of Source Address Validation in the Internet. In: CCS'19. ACM (2019)
33. Mauch, J.: Spoofing ASNs, http://seclists.org/nanog/2013/Aug/132
34. MaxMind LLC: Maxmind geoip, https://www.maxmind.com/en/geoip2-databases
35. McConachie, A.: Anti-Spoofing, BCP 38, and the Tragedy of the Commons, https://www.internetsociety.org/blog/2014/07/anti-spoofing-bcp-38-and-the-tragedy-of-the-commons/
36. Müller, L.F., Luckie, M.J., Huffaker, B., kc claffy, Barcellos, M.P.: Challenges in inferring spoofed traffic at IXPs. In: ACM Conference on Emerging Networking Experiments And Technologies (CoNEXT). pp. 96–109 (2019)
37. Noroozian, A., Ciere, M., Korczyński, M., Tajalizadehkhoob, S., van Eeten, M.: Inferring security performance of providers from noisy and heterogenous abuse datasets. In: WEIS (2017)
38. Padmanabhan, R., Dhamdhere, A., Aben, E., claffy, k., Spring, N.: Reasons Dynamic Addresses Change. In: Internet Measurement Conference. ACM (Nov 2016)
39. Rapid7: https://opendata.rapid7.com/sonar.udp/
40. Rossow, C.: Amplification Hell: Revisiting Network Protocols for DDoS Abuse. In: NDSS (2014)
41. Senie, D., Ferguson, P.: Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing. RFC 2827 (May 2000), https://rfc-editor.org/rfc/rfc2827.txt
42. Skwarek, M., Korczynski, M., Mazurczyk, W., Duda, A.: Characterizing vulnerability of DNS AXFR transfers with global-scale scanning. In: 2019 IEEE Security and Privacy Workshops. pp. 193–198 (2019)
43. SpamHaus: https://www.spamhaus.org/
44. Tajalizadehkhoob, S., Böhme, R., Gañán, C., Korczyński, M., van Eeten, M.: Rotten Apples or Bad Harvest? What We Are Measuring When We Are Measuring Abuse. ACM Trans. Internet Techn. **18**(4), 49:1–49:25 (2018)
45. TeleGeography: https://www.telegeography.com/products/globalcomms/
46. Van Eeten, M., Bauer, J.M., Asghari, H., Tabatabaie, S., Rand, D.: The role of internet service providers in botnet mitigation an empirical analysis based on spam data. TPRC (2010)
47. Vixie, P.: Rate-limiting state. Communications of the ACM **57**(4), 40–43 (2014)
48. Zhang, J., Durumeric, Z., Bailey, M., Liu, M., Karir, M.: On the mismanagement and maliciousness of networks. In: Network and Distributed System Security Symposium (NDSS). The Internet Society (2014)