Aliens Among Us: Observing Private or Reserved IPs on the Public Internet

Radu Anghel
TU Delft
r.anghel@tudelft.nl

Carlos Gañán
ICANN
carlos.ganan@icann.org

Qasim Lone
RIPE NCC
qlone@ripe.net

Matthew Luckie
CAIDA
mjl@caida.org

Yury Zhauniarovich
TU Delft
y.zhauniarovich@tudelft.nl

***Abstract*—Spoofed traffic remains a major network hygiene concern, as it enables Distributed Denial-of-Service (DDoS) attacks by obscuring attack origins and hindering forensic analysis. A key indicator of poor hygiene is the presence of *Bogon* traffic—packets carrying invalid or non-routable source addresses—in the public Internet, arising from misconfigurations or insufficient filtering. Despite long-standing Source Address Validation (SAV) recommendations such as BCP 38 and BCP 84, Bogon filtering remains inconsistently deployed. In this work, we analyze eight years (2017–2024) of traceroute measurements from the CAIDA Ark platform, enriched with historical BGP data from RIPE RIS and RouteViews, to quantify the prevalence and characteristics of Bogon addresses in the data plane. We observe widespread non-compliance with best practices: between 82.69% and 97.83% of Ark vantage points encounter traceroute paths containing Bogon IPs, predominantly RFC1918 addresses. Overall, 21.11% of traceroutes include RFC1918 addresses, with smaller fractions involving RFC6598 (1.68%) and RFC3927 (0.08%). We identify over 15,500 Autonomous Systems (ASes) that transit Bogon traffic, although only 11.88% do so in more than half of the measurements. Cross-referencing with the Spoofer project and MANRS reveals a significant gap between control- and data-plane assurances: 52.71% of ASes forwarding Bogon-sourced packets are classified as non-spoofable, indicating incomplete or ineffective SAV deployment.**

I. Introduction

RFC3871 [1] defines *Bogons* as packets with source IP addresses from address blocks that are either unallocated by IANA or Regional Internet Registries (RIRs), or reserved for private or special use. These packets originating from addresses that do not have valid return routes in the Global Routing Table and, thus, are not globally reachable, should never appear on the public Internet. Their presence can be considered as one manifestation of spoofed traffic.

For many years, spoofed Internet traffic has been regarded as a symptom of poor network hygiene. Such traffic enables Distributed Denial-of-Service (DDoS) attacks, which generate large volumes of data either directly or via amplification mechanisms [2], [3], [4]. By obscuring the true origin of attacks, spoofing significantly hinders forensic analysis and often makes attribution infeasible. Consequently, mitigating spoofed traffic has become a long-standing priority within the network security community [5], [6], [7]. This effort has led to the development of a class of technical defenses known collectively as Source Address Validation (SAV) [8], [9], which aim to verify packet source addresses and reduce the impact of IP spoofing.

Despite growing awareness of cybersecurity principles like security by design and secure defaults [10], these concepts are not yet widely implemented in practice. As a result, network devices often require manual configuration adjustments to improve their security. Guidance documents such as BCP 38 [11] and BCP 84 [12], [13] provide recommendations for implementing SAV on both ingress and egress traffic. Researchers have developed several approaches to assess SAV deployment, including active measurements originating from the tested network [14], traceroute loop analysis [15], and sending spoofed packets to DNS resolvers [16]. Meanwhile, initiatives like MANRS [17] seek to incentivize network operators to adopt these practices. Yet, Bogon packets are frequently overlooked and dismissed as a minor issue, rather than being treated as a substantial security problem.

This paper challenges the assumption that Bogons are inconsequential. If an AS transits Bogons, it means it does not implement proper ingress or egress filtering for transit traffic. As a result, it can be exploited to carry spoofed traffic, forcing both the AS and its upstream providers to spend resources on transporting and processing unnecessary packets. Such traffic can lead to false positives in security systems, which commonly flag Bogon packets as suspicious. If a Bogon address is used inside an AS and is visible outside, this means that this AS does not perform filtering and thus, can be a source of spoofed traffic. This information might help in reconnaissance, e.g., by discovering which Private-Use (RFC1918) addresses are used by the network [18], [19]. This information, combined with the lack of filtering and additional knowledge (i.e., SNMP communities), could allow attackers to exploit protocols that work over UDP, such as SNMP, by spoofing a potentially trusted source to trigger resets or configuration changes on Customer Premises Equipment [20].

In this study, we investigate the prevalence of ASes that likely do not adhere to SAV best practices by examining the presence of Bogon addresses in traceroute measurements collected by CAIDA Ark. By analyzing data spanning eight

Network and Distributed System Security (NDSS) Symposium 2026
23-27 February 2026, San Diego, CA, USA
ISBN 979-8-9919276-8-0
https://dx.doi.org/10.14722/ndss.2026.241118
www.ndss-symposium.org

years from January 2017 to December 2024 and applying our own heuristics, we identify ASes that most likely deviate from established SAV best practices. Our primary motivation was to introduce an alternative approach for identifying ASes that transit or originate spoofed traffic, addressing limitations in existing efforts like Spoofer and MANRS, which rely on voluntary AS collaboration.

Our analysis reveals the presence of Bogon addresses and a lack of filtering beyond the AS borders within thousands of ASes. Between 82.69% and 97.83% of CAIDA Ark vantage points detect at least one path containing Bogon IP addresses, with the most prevalent being from the RFC1918 [21] space. We find that 21.11% of the analyzed traceroutes contain RFC1918 [21] addresses, 1.68% – RFC6598 [22], and 0.08% include addresses from RFC3927 [23]. We identified 15,541 unique ASes transiting Bogons. While some ASes traverse multiple Bogon types, ASes transiting RFC1918 [21] addresses dominate. However, we observe that ASes transiting Bogons do not appear consistently across many measurements. The total number of ASes that appeared transiting Bogons of all types across all 96 measurements is 283. Only 11.88% (1,846) of ASes are observed in more than half of the measurements. This behavior likely results from either measurement inconsistencies in the Ark dataset due to the randomization of vantage points and destinations or ASes fixing their misconfigurations.

We compared our results with two other sources tackling SAV: the Spoofer project data [14] and the Mutually Agreed Norms for Routing Security (MANRS) [17]. Cross-checking our findings with Spoofer (considering measurements conducted within 6 months before the corresponding traceroute with an AS transiting packets with Bogon sources, as explained in Section III), we found that 52.71% of the unique ASes transiting RFC1918 [21] Bogons that are also present in the Spoofer dataset (2,529) were reported as non-spoofable.

Thus, the main contributions of this paper are:

- We present the first comprehensive assessment of network hygiene with a focus on packets with Bogon source addresses moving beyond conventional concerns strictly related to IP spoofing. Our findings reveal that private addresses are more prevalent than other types of special-purpose Bogons appearing on the Internet.
- We identify widespread non-compliance with best practices across various ASes, mostly in ISP networks. Over 56% of these ASes are registered in the top ten countries, with the largest shares located in the USA, Brazil, and Russia.
- We provide a detailed breakdown of ASes transiting multiple types of Bogons, offering insights into their distribution and patterns. Over the study period, 15,541 unique ASNs were found transiting packets sourced by Bogon addresses.
- We characterize the intersection between spoofable and non-spoofable ASes, as found by the Spoofer project, revealing an overlap with ASes forwarding packets with Bogons as source address. We provide recommendations to address the gap in the deployment of SAV measures.

II. Background

A. Bogon Addresses

Bogons refer to packets with source addresses that are not supposed to be routed on the public Internet. These include packets originating from unallocated ranges, i.e., IP addresses reserved but not yet assigned by a Regional Internet Registry, as well as from special-use ranges, some of which are defined in the IANA IPv4 Special-Purpose Address Registry [24]. Since unallocated ranges change over time, in this study, we focus specifically on the following Special-Purpose ranges:

240.0.0.0/4: As per RFC1112 [25], these "Class E" IP addresses are reserved for future use and not publicly routable, though recent studies show their internal use by large providers such as Amazon AWS [26].

127.0.0.0/8: These IPs are used as Loopback IPs that always point to the local host (RFC1122 [27]).

10.0.0.0/8, 172.16.0.0/12, 192.168.0.0/16: Per RFC1918 [21], these blocks are Private-Use addresses widely used behind firewalls and NATs.

169.254.0.0/16: According to RFC3927 [23], these Link-Local addresses are used to permit IP connectivity between hosts in the same physical network when no static or dynamic IP configuration is provided to them.

192.0.2.0/24, 198.51.100.0/24, 203.0.113.0/24: These addresses, specified in RFC5737 [28], are reserved for use in documentation, specifications and examples.

100.64.0.0/10: According to RFC6598 [22], this block of IP addresses is anticipated to be used as Shared Address Space for Carrier-Grade NAT (CGN) devices. Their purpose is similar to Private-Use (RFC1918) addresses, but they are intended to be used on Service Provider networks.

192.0.0.0/24: These IPs, defined in RFC6890 [29], are called Protocol Assignments and reserved for use by various protocols, such as NAT64/DNS64 Discovery (RFC8880 [30] / RFC7050 [31]).

192.88.99.0/24: According to RFC7526 [32], this block of 6to4 Relay Anycast addresses was initially designed to aid transition to IPv6, now deprecated due to high failure rates.

B. Implications of Unfiltered Bogon Addresses

Failure to filter packets originating from Bogon addresses can enable a range of attack scenarios and security risks. Here, we outline some of the potential threats.

SNMP-Enabled Exploitation. Access Control Lists (ACLs) of SNMP agents are often configured to trust internal network segments. If spoofed packets claiming trusted private addresses (for example, 192.168.0.0/16 or 10.0.0.0/8) are not filtered, and the SNMP `write community`[1] is known (default value is equal to `private`), an adversary can send SNMP SET requests to vulnerable devices. Such requests can enable various malicious actions, including forced reboots, unauthorized configuration changes, and even full compromise of the device [20]. Exposed devices with read-only access

[1] *Write community* is a string protecting SNMP device write operations.

could also be fully compromised via exploit chains [33], [34], [35]. Moreover, SNMP-enabled devices can be abused to amplify denial-of-service attacks [36]: the protocol can reach the amplification factor of up to 650 [37], [38].

Network Reconnaissance and Topology Discovery. Bogon addresses appearing on the public Internet can leak internal addressing schemes and enable remote network mapping [18], [19]. This form of exposure corresponds to techniques described in the MITRE ATT&CK framework under System Network Configuration Discovery (T1016) [39], which explains how adversaries use network configuration details to identify potential paths to access target networks. Public visibility of private addressing also eases automated discovery workflows that locate vulnerable infrastructure segments.

Spoofing Attacks and Attribution Evasion. Bogon filtering failures increase Internet-wide exposure to spoofing attacks [40]. Source IP address spoofing allows attackers to hide their identity and impersonate legitimate networks [41]. This can mislead forensic analysis into believing attacks originated from innocent third parties, potentially causing collateral damage when defenders implement reactive blocking of spoofed source networks [11], [42]. Attribution of such attacks is exceedingly difficult when packets with Bogon sources traverse multiple ASes to reach their destination.

Infrastructure and Operational Impact. In addition to the direct security risks, transiting packets with Bogon origins imposes operational burdens on all networks on the path, not only the source and destination networks, since they need to process the packets, leading to wasted bandwidth and increased computational overhead on the routing infrastructure [43]. Security monitoring systems can also experience elevated false-positive rates due to the presence of Bogon packets [44], which can create alert fatigue and potentially mask genuine threats. Moreover, there is other specific operational impact related to Private-Use (RFC1918) addresses [45].

C. Datasets

In this work, we rely on several datasets, which collectively provide valuable insights on networks engaged in the Bogon filtering practices and the deployment of SAV mechanisms.

CAIDA IPv4 Routed /24 Topology. For our research, we use snapshots of the CAIDA IPv4 Routed /24 Topology Dataset [46], collected through the Ark measurement infrastructure [47]. It comprises globally distributed monitors (Vantage Points, VPs) that use scamper [48] to probe a random IP address in each routed IPv4 /24 prefix on a daily cycle. We employ the resulting traceroute measurements, referring to them as the *Ark Dataset*.

RIPE RIS and UOregon RVIEWS. To identify ASes that transit packets with Bogon source addresses, we use BGP routing information snapshots. The RIPE NCC's Routing Information Service (RIS) project [49] and the University of Oregon's Route Views project [50] both collect and archive global BGP routing information. RIS currently operates 26 route collectors (`RRCxx`) at Internet Exchange Points (IXPs) across five continents. Three of them operate as "multi-hop" collectors, meaning that they collect routes not only from networks present at that particular IXP but also from the rest of the world. The University of Oregon's Route Views project [50] also collects global routing information from 42 such collectors distributed worldwide. Both projects provide data in Multi-threaded Routing Toolkit (`MRT`) format (RFC6396 [51]), available as full dumps and incremental updates. Further, we refer to these datasets as the *RIPE RIS* and *UOregon RVIEWS* datasets, respectively.

CAIDA Spoofer. The Spoofer project [14] aims to evaluate and report on the implementation of best practices for source address validation (SAV) to mitigate spoofing attacks. The project uses client and server software to send and receive packets with spoofed source addresses. Networks that fail to drop invalid packets, either inbound from internal addresses or outbound to external addresses, are considered vulnerable. Test results with anonymized IPs are publicly available, while full results with real IPs are shared privately with network operators to help identify missing SAV deployment. We employ the publicly available dataset to validate our findings.

MANRS for Network Operators. The *Mutually Agreed Norms for Routing Security (MANRS)* [17] initiative, led by Internet Society, aims to foster collaboration among network operators, IXPs, equipment vendors, CDN and cloud providers to enhance the resilience and security of the routing infrastructure. By joining the initiative, they show their commitment to achieving this goal, while the initiative tracks participants' commitment by performing a set of regular measurements, showing their readiness to adhere to *Compulsory* and *Recommended* actions stemming from recognized industry best practices. The results of these measurements are publicly available. One group of MANRS participants, *Network Operators* [52], is encouraged to implement the `anti_spoofing` action, which checks whether packets with spoofed IP addresses can be sent through their networks. We employ the results of this check in our validation study.

III. Methodology

To identify ASes permitting Bogons transit, we employ the Ark dataset, containing the route information between a source, represented by an Ark Vantage Point (*Ark VP*), and a destination, a randomly selected IP address from each routed IPv4 /24 prefix. Ark VPs collect this information using the traceroute tool, which obtains traceroute paths by sending a series of packets, typically *ICMP Echo Requests*, to the same destination, incrementally increasing Time-to-Live (TTL) values. These values are used to prevent packets from circulating indefinitely: each router along the path decreases TTL by 1. When a router receives a packet with TTL of 1, it discards the packet and sends back an *ICMP Time Exceeded* error message, with the router's IP address as the source. By interpreting these responses, the traceroute tool reconstructs the route a packet takes to reach the destination.

Because the traceroute tool records the Source IP of the incoming ICMP Time Exceeded message, any AS on the

return path from the router to the VP is effectively transiting a packet with a Bogon source address.

The rationale behind our approach is as follows. If a collected traceroute path in the Ark dataset contains a Bogon IP address, it implies that an *ICMP Time Exceeded* error message successfully reached the source and was sent from a router configured with this Bogon address on an interface or loopback. Because the traceroute tool records the source IP of the incoming ICMP Time Exceeded message, any AS on the return path from the router to the VP is effectively transiting a packet with a Bogon source address. However, according to best practices, packets originating from Bogon IP addresses should be filtered and should not be visible on the public Internet. Therefore, the presence of such addresses in traceroute paths indicates that routers along the path (and, by extension, the corresponding ASes) fail to implement proper filtering, suggesting non-compliance with best practices and potentially SAV. For instance, consider the example shown in Figure 1, where the Ark VP traces the route to the destination (*DST*). The collected path – RN $\rightarrow$ ... $\rightarrow$ R6 $\rightarrow$ R5 $\rightarrow$ R4 $\rightarrow$ R3 $\rightarrow$ R2 $\rightarrow$ **RB** $\rightarrow$ R1 $\rightarrow$ R0 – includes the IP addresses of the routers encountered along the way, with **RB** assigned a Bogon (i.e., invalid) IP address. This strongly suggests that router R2 (and, by extension, AS64500) does not perform proper egress filtering, while router R3 (and AS65550) fails to implement ingress filtering.

Figure 2 illustrates the methodology to identify and characterize ASes transiting Bogons. The rectangles with numbered circles indicate specific methodology steps, while document icons represent the data used in each step.

Step ①: Collect Ark traceroutes. At this step, we extract traceroutes from the Ark dataset [46]. Due to the large volume of traceroute data collected by CAIDA Ark, analyzing every daily snapshot would be computationally prohibitive. To balance computational feasibility with adequate temporal coverage, we selected one representative 18^{th} day[2] per month over a 8-year period, from January 2017 to December 2024. When multiple measurement cycles occur the same day, we select traceroutes from the cycle with the lowest number.

Step ②: Identify traceroute hops with Bogon addresses. At this step, we check if an IPv4 address of a traceroute hop belongs to one of the selected IANA IPv4 Special-Purpose Addresses, listed in Table I and described in Section II-A. We do not search for addresses coming from prefixes such as 0.0.0.0/8 [53], [27], or 192.31.196.0/24 (AS112-v4) [54], and 192.175.48.0/24 (Direct Delegation AS112 Service) [55] due to the low probability of these addresses being visible as intermediary hops in traceroutes.

[2]Initially, we chose the 1^{st} day of each month; however, we found that several months contained incomplete cycles (with very few traceroutes), or all data for the whole day was missing. We sequentially examined several next days, specifically, between 1^{st} and 12^{th}, but encountered similar data gaps. Therefore, we started checking randomly other days and discovered that the 18^{th} day of each month consistently provided complete and continuous data across the entire observation period.

TABLE I: Bogon address blocks

RFC	Description	CIDR
1112	Former Class E	240.0.0.0/4
1122	Loopback	127.0.0.0/8
1918	Private-Use	10.0.0.0/8, 172.16.0.0/12, 192.168.0.0/16
3927	Link-Local	169.254.0.0/16
5737	Documentation	192.0.2.0/24, 198.51.100.0/24, 203.0.113.0/24
6598	Shared Address Space	100.64.0.0/10
6890	Protocol Assignments	192.0.0.0/24
7526	6to4 Relay Anycast	192.88.99.0/24

Step ③: Add originating AS information to traceroute hops. During this step, we map each hop's IPv4 address in a traceroute to the corresponding origin ASN. To perform this task, we use two datasets containing historical routing information, namely RIPE RIS and UOregon RVIEWS (see Section II-C). For this, we download the `MRT` dump files collected at *00:00* on the same day as the Ark traceroute cycle. From the RIPE RIS dataset, we used RIBs gathered by the `RRC00` collector. This collector stands out because it is one of the multi-hop RIS route collectors that consolidates information from a wide array of peers found in various locations globally. From the UOregon RVIEWS dataset, we use the dumps from the `route-views2` collector. In some cases, RIPE RIS and UOregon RVIEWS datasets also have a mapping between some special-purpose IP addresses and a corresponding AS. These are prefixes leaked to route collectors, and we ignore the origin AS found for them in these datasets. Based on this information, we build an *AS path* for the data plane of each traceroute – a sequence of ASNs, each of which corresponds to a hop in the traceroute.

Step ④: Clean the AS path. We remove consecutive duplicate ASNs (multiple hops in the same ASN), hops with *unknown* origin that are not Bogon addresses (public addresses not found in the Global Routing Table, usually IXP prefixes).

Step ⑤: Categorization. After obtaining a clean AS path, we look for Bogon addresses that are found at more than one AS-hop away from the source of the traceroute, ignoring the Bogons inside the network originating the traceroute. We do this since it is expected to have private addresses inside your own network; however, they should not be visible when 'crossing the border' to a different AS. We specifically look for the following cases:

BA: All ASes on the path before a Bogon address. Since the packets with Bogon addresses as source reach the origin of the traceroute (Ark VP), this case shows that all ASes on that path forward these packets, against best practices. It is, however, possible that due to the asymmetry of routing on the Internet, the actual reverse path these packets take to the origin of the traceroute could be different from the forward path. Thus, there is less certainty about the filtering practices of these ASes, except for the one closest to the origin of the traceroute. In Figure 1, AS64500, AS65550, and AS65540 belong to this category.

BB: The AS found right before a Bogon address. The

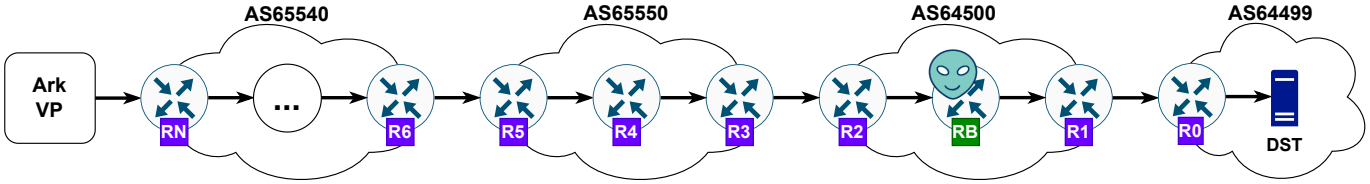

Fig. 1: Approach Idea

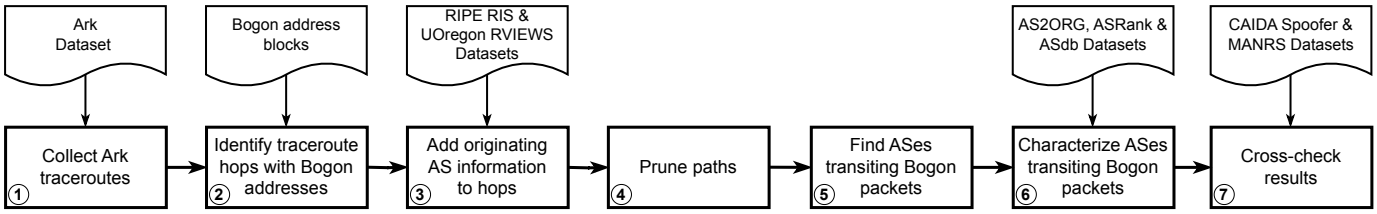

Fig. 2: Bogon Addresses Identification Overview

AS originating the address found on the traceroute right before a Bogon address could mean that Bogon addresses are in use inside that AS, or a neighboring AS is using them. In both cases though the AS found immediately before the Bogon address is transiting those packets with a Bogon address as a source and forwarding them beyond their AS towards the Ark VP running the traceroute. In Figure 1, AS64500 belongs to this category.

BC: An AS sandwich with a Bogon inside. If the same AS is found both before and after Bogon addresses, this signals that, with a high probability, the Bogon addresses are in use inside this AS. Due to BGP loop prevention mechanisms, it is highly unlikely that packets would leave the AS preceding the Bogon address, traverse to a different AS, and then return to the original AS for the traceroute hop following the Bogon. That AS also does not filter the corresponding packets as they are forwarded towards the Ark VP. In Figure 1, AS64500 belongs to this category.

Note that according to this definition BC is a subset of BB, which is a subset of BA ($BC \subseteq BB \subseteq BA$). Thus, in our work, we consider two situations when traceroutes uncover ASes either 1) actively using Bogon addresses (BC and BB cases) or 2) transiting Bogon packets (BB and BA cases). The BC ASes use Bogon addresses inside their networks, the BB ASes could be both containing and transiting, and BA are mostly transiting (assuming no multipath).

Step ⑥: Characterize ASes transiting Bogons. During this step, we characterize the ASes identified in previous steps by using publicly available datasets: ASdb [56] (as of 2024-01), AS2ORG [57] (as of 2025-01-01) and ASRank [58] (as of 2025-01-01). The ASdb dataset [56] categorizes organizations associated with an AS according to the North American Industry Classification System (NAICSlite). For our characterization, we use the 'Category 1 - Layer 1' and 'Category 1 - Layer 2' fields. We employ the AS2ORG dataset [57] to provide a mapping between the AS and the Regional Internet Registry (RIR) where the AS is registered, and identify the country of the organization. Finally, we use the ASRank [58] dataset to obtain the geographical coordinates of the registration place of the AS in order to build the map graph.

Step ⑦: Cross-check the results. Here we compare our findings with the results in two other relevant datasets: CAIDA Spoofer [14] and MANRS for Network Operators [52].

Transiting packets with Bogon addresses as source or destination is a clear sign of a misconfiguration, but it does not necessarily mean that IP spoofing is also possible on that network. For this, the Spoofer dataset acts as ground truth due to the active measurements performed from within the tested AS. We consider an AS as spoofable if we can find at least one match in the Spoofer dataset during a period of up to 6 months before the corresponding Ark traceroute measurement. We use such a large timespan of Spoofer data because older versions of the Spoofer client required checks to be manually started, i.e., the end-user has to run the software and push a button to start the test, the result of this approach being irregular measurements. Recent versions of the Spoofer client perform the tests automatically and run the measurements regularly. In order to match with the period of Ark traceroutes, we downloaded through the Spoofer API results for tests run between January 2017 and December 2024.

Although MANRS `anti-spoofing` action checks rely on data from the Spoofer project, the added information is that the network operator actively pledged to undertake anti-spoofing measures. For MANRS, we downloaded the latest list of conformance for all network operators participating in MANRS (as of 2025-07-22) by using the MANRS API,[3] we then cross checked that data with the set of ASes we identified as transiting packets with Bogon source or destination.

Note that the Spoofer measurements are done on a voluntary basis and thus, are not performed from all ASes. Also, participants can opt out of sharing the results publicly. Thus, the Spoofer dataset may be incomplete.

[3]https://manrs.stoplight.io/docs/manrs-public-api/97379961794c7-list-net-ops-conformance

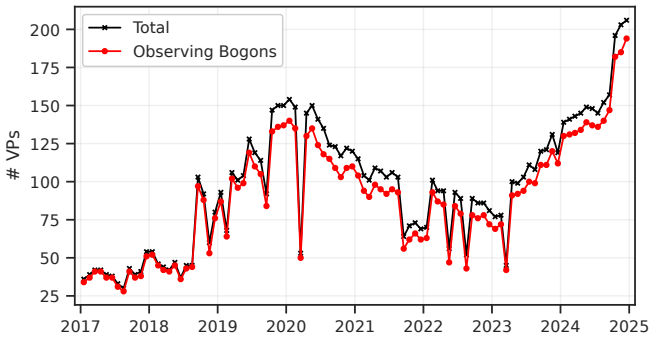

Fig. 3: Number of Vantage Points

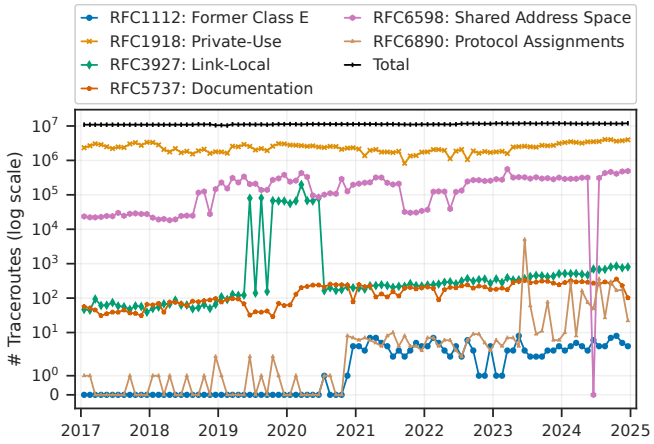

Fig. 4: Number of Traceroutes with Bogon addresses

IV. Findings

To identify the existence and prevalence of Bogon addresses on the public Internet, we conducted a longitudinal study of traceroutes collected by the CAIDA Ark project over an eight-year period. We downloaded the traceroute data for the 18^{th} day of each month within this period and analyzed them following the methodology described in Section III. Table II reports on the data used in this study and summarizes our findings. In this section, we present the main results of our analysis: (A) prevalence of Bogon addresses in traceroutes, (B) ASes transiting Bogons, and (C) characteristics of those ASes.

A. Traceroutes with Bogon Addresses

Figure 3 shows the number of Ark Vantage Points (VPs) used to collect traceroutes. The black line with cross markers shows the total number of Vantage Points used to perform the measurements (the "# VPs" column in Table II reports the same numbers). As we can see, the number of VPs was pretty low (under 40 VPs) till August 2018, then started to grow, reaching 154 VPs in January 2020, and afterward, decreasing till March 2023. In general, the number of VPs is quite unstable (97.2 ± 41.6), with deep drops happening from time to time (e.g., in March 2020).

The red line with dot markers shows the number of VPs observing at least one Bogon address (the "# VPs observing Bogons" column in Table II). Between 82.69% and 97.83% of CAIDA Ark VPs observed at least one traceroute path containing a Bogon address during a measurement.

The black line with vertical line markers in Figure 4 shows the total number of analyzed traceroutes per month. We analyzed 1,085,410,451 traceroutes in total. The average number of traceroutes per measurement is 11,240,248, with a maximum of 11,998,993 recorded in December 2024 and a minimum of 10,330,500 observed in February 2019.

The other lines in Figure 4 correspond to the number of traceroutes containing Bogon addresses of particular types defined in Table I. As we can see, the most common traceroutes with Bogon IPs contain Private-Use (RFC1918) addresses. They are found in 229,166,687 (21.11%) of all traceroutes, with an average of 2,387,152 per measurement. Most often, these Private-Use (RFC1918) addresses are located close to the traceroute source and do not cross the AS border. We believe they are used in the infrastructure of the ASes hosting the Ark nodes. Protocol Assignments (RFC6890) addresses were observed in 7,020 traceroutes. By analyzing the paths, we believe that these addresses are used similarly to the Private-Use (RFC1918) addresses and not for the specific protocols they were reserved for. We make this assumption because we observe IPs such as 192.0.0.201, a part of 192.0.0.0/24 prefix that is not assigned to any specific protocol. Most of the observed paths lead to the ASN of a telecom provider in Brazil, to an ASN belonging to a cloud service provider in the US, and to four ASNs located in Spain. The Former Class E (RFC1112) addresses appear in 189 traceroutes over the observed period. All these traces target only two ASNs, the largest AS in Japan, according to IHR [59], and one large US-based AS. Although it is known that networks such as Amazon AWS use these addresses internally [26], they were not visible in the Ark dataset, suggesting proper filtering is in place. Documentation (RFC5737) prefixes are found in 15,377 traceroutes. From 1,085,410,451 analyzed traceroutes, 1.68% contain Shared Address Space (RFC6598) IPs, and 0.08% include Link-Local (RFC3927) addresses. Interestingly, there is a spike in the number of traces with Link-Local (RFC3927) addresses between June 2019 and July 2020. Moreover, at the start of this period, namely in July and September, the number of traceroutes with this type of Bogon addresses returned to their previous values. These spikes are likely due to temporary misconfigurations inside these ASN(s), which were fixed later.

Other types of Bogon addresses are found in less than 0.01% of the traceroutes. Loopback (RFC1122) and the 6to4 Relay Anycast (RFC7526) addresses do not appear in any traceroutes, therefore, we exclude them from further consideration.

B. ASNs Transiting Packets with Bogon Addresses

In this section, we report the results of the analysis of ASNs transiting Bogons as described in Section III (see Step ⑤). Figure 5 presents the breakdown of ASes transiting Bogons

TABLE II: Dataset Statistic

Month	# VPs	# VPs observing Bogon addresses	# Traceroutes								# ASNs transiting Bogons per RFC					
			Total	With Bogon addresses per RFC						W/ no ASNs						
				1112	1918	3927	5737	6598	6890		1112	1918	3927	5737	6598	6890
2017-01	36	34	10,908,012	0	2,333,712	47	57	23,535	1	229,470	0	2,565	61	39	251	4
2017-02	39	37	10,908,012	0	2,665,121	45	51	22,179	1	235,726	0	2,508	64	43	261	4
2017-03	42	41	10,908,012	0	3,032,516	94	45	22,113	0	242,343	0	2,512	83	41	255	0
2017-04	42	41	10,908,012	0	2,864,580	61	31	22,769	0	265,457	0	2,484	69	32	267	0
2017-05	39	37	10,903,071	0	2,496,062	62	35	24,231	0	267,989	0	2,517	69	49	268	0
2017-06	38	37	10,908,012	0	2,211,989	74	39	23,915	0	271,619	0	2,568	76	35	289	0
2017-07	33	31	10,908,012	0	2,449,291	59	39	29,861	0	282,081	0	2,416	57	40	264	0
2017-08	30	28	10,908,012	0	2,369,351	56	45	24,570	0	288,012	0	2,472	49	32	252	0
2017-09	43	41	10,906,323	0	3,009,100	47	37	28,063	1	291,330	0	2,352	64	47	269	4
2017-10	39	37	10,908,012	0	3,275,053	59	36	28,838	0	294,498	0	2,253	71	40	278	0
2017-11	41	38	10,906,398	0	2,753,270	57	31	27,807	0	293,535	0	2,287	63	29	267	0
2017-12	54	51	10,906,262	0	3,362,171	39	66	27,597	0	299,026	0	2,190	61	45	274	0
2018-01	54	52	10,906,280	0	3,323,094	50	63	21,706	1	319,112	0	2,408	71	43	293	4
2018-02	46	45	10,904,557	0	2,829,771	55	70	19,225	1	328,951	0	2,138	70	33	262	4
2018-03	44	42	10,908,012	0	2,075,940	68	39	19,851	0	347,956	0	2,186	67	31	259	0
2018-04	42	41	10,906,242	0	1,766,926	67	76	18,263	1	333,916	0	2,169	72	39	277	4
2018-05	47	45	10,906,404	0	2,209,012	87	76	19,310	0	353,910	0	2,299	72	44	265	0
2018-06	37	36	10,905,421	0	1,680,677	63	70	24,331	0	372,640	0	2,158	58	35	254	0
2018-07	45	43	10,908,012	0	1,841,169	63	60	24,849	1	377,478	0	2,281	73	38	281	4
2018-08	45	44	10,906,464	0	1,538,555	49	81	24,750	0	393,395	0	1,914	75	42	293	0
2018-09	103	97	11,194,865	0	1,891,850	53	78	115,758	0	144,702	0	2,091	81	63	389	0
2018-10	92	88	11,075,865	0	2,106,341	65	84	124,939	1	156,946	0	2,495	79	46	370	2
2018-11	60	53	11,204,865	0	1,608,819	50	88	27,600	0	315,055	0	2,595	64	36	359	0
2018-12	80	76	10,466,865	0	1,777,668	66	97	145,832	2	394,152	0	2,735	79	55	350	2
2019-01	93	87	10,547,865	0	1,760,943	109	77	227,752	1	399,736	0	2,479	100	38	363	4
2019-02	68	64	10,330,500	0	1,621,149	91	92	153,034	0	399,084	0	2,532	71	45	290	0
2019-03	106	102	11,204,730	0	2,578,371	128	96	309,956	0	439,286	0	2,760	123	62	410	0
2019-04	101	96	10,999,256	0	2,489,626	120	94	228,744	0	77,086	0	2,796	113	59	436	0
2019-05	104	99	11,107,221	0	2,907,541	122	68	338,754	0	98,394	0	2,766	124	46	439	0
2019-06	128	119	11,126,128	0	2,557,159	79,620	32	204,499	2	147,756	0	2,676	106	46	437	2
2019-07	119	110	11,031,628	0	2,010,701	140	40	210,255	0	149,073	0	2,683	119	52	461	0
2019-08	114	105	11,118,628	0	2,203,410	81,142	39	138,174	1	162,203	0	2,732	121	45	497	2
2019-09	92	84	10,911,628	0	1,933,021	155	42	138,969	0	182,851	0	2,759	114	33	464	0
2019-10	147	133	11,119,216	0	2,588,404	67,630	29	271,362	2	204,715	0	2,551	113	36	480	6
2019-11	150	136	11,308,704	0	3,066,096	66,138	70	299,519	1	368,798	0	2,682	120	66	488	3
2019-12	150	137	11,269,704	0	3,015,695	66,128	61	380,526	0	398,909	0	2,551	116	49	477	0
2020-01	154	140	11,331,204	0	2,795,546	55,620	63	241,246	0	412,223	0	2,521	114	55	504	0
2020-02	149	135	11,259,000	0	2,760,966	66,159	129	258,399	1	419,925	0	2,439	135	67	477	1
2020-03	53	50	11,118,985	0	2,676,252	198,138	202	429,860	0	429,162	0	2,173	84	46	380	0
2020-04	145	130	11,332,908	0	2,568,605	67,630	219	330,081	0	440,969	0	2,640	116	82	486	0
2020-05	150	135	11,331,408	0	2,635,751	66,149	238	95,954	0	504,577	0	2,692	130	88	279	0
2020-06	141	124	11,334,204	0	2,446,309	79,636	239	86,206	0	506,685	0	2,609	131	95	274	0
2020-07	135	118	11,325,000	1	2,375,570	165	210	101,066	1	496,324	3	2,637	125	85	264	4
2020-08	124	115	11,326,500	0	2,551,634	201	254	111,035	1	496,130	0	2,675	146	94	276	3
2020-09	123	109	11,286,204	0	2,520,081	167	243	107,487	0	498,144	0	2,683	124	83	276	0
2020-10	117	103	11,326,500	0	2,106,997	176	243	289,132	0	514,640	0	2,757	143	75	524	0
2020-11	122	109	11,290,908	1	2,273,141	216	238	126,485	8	510,543	2	2,701	148	115	544	17
2020-12	120	110	11,328,204	4	2,324,084	200	78	196,482	7	235,925	11	2,642	151	69	540	23
2021-01	115	104	11,329,704	4	2,140,406	198	248	210,825	6	528,196	6	2,587	137	95	524	18
2021-02	104	94	11,334,914	3	1,369,941	183	213	224,061	7	526,893	8	2,363	147	84	522	12
2021-03	101	90	11,334,408	7	1,957,880	228	226	229,972	6	527,306	13	2,599	159	82	555	22
2021-04	109	98	11,329,500	7	2,064,681	231	108	318,913	5	557,966	8	2,648	162	80	549	10
2021-05	107	95	11,200,500	5	1,751,453	243	132	317,190	4	531,416	6	2,599	161	79	511	12
2021-06	103	92	11,332,908	4	1,754,968	231	108	223,049	8	541,410	9	2,504	130	64	525	17
2021-07	106	95	11,314,500	2	1,694,122	205	154	200,085	10	553,234	5	2,471	139	68	487	19
2021-08	103	93	11,193,204	3	1,836,343	220	115	210,469	4	549,035	3	2,486	132	86	499	5
2021-09	64	56	11,109,408	2	822,262	220	195	31,948	8	554,082	6	2,531	127	79	514	17
2021-10	71	62	11,028,204	3	1,343,907	263	205	30,237	4	550,561	9	2,504	141	85	507	11
2021-11	73	66	11,106,000	5	1,383,787	230	185	30,622	4	561,695	10	2,437	124	81	519	13
2021-12	69	62	11,172,000	4	1,739,596	224	200	34,095	3	566,579	10	2,539	130	88	534	11
2022-01	70	63	11,169,000	4	1,769,730	240	212	36,863	7	570,746	9	2,513	127	78	529	11
2022-02	101	93	11,208,204	7	2,061,262	245	190	123,164	7	578,420	12	2,728	141	93	530	14
2022-03	94	87	11,205,204	5	2,068,627	256	89	126,164	4	583,457	6	2,594	128	79	523	10
2022-04	94	85	11,208,204	3	1,956,568	288	176	124,553	6	590,618	7	2,545	136	67	543	11
2022-05	56	47	11,073,408	2	1,122,791	297	206	38,957	6	585,876	4	2,705	145	63	538	14
2022-06	93	84	11,209,502	4	1,866,069	259	200	121,859	3	609,576	6	2,567	144	69	536	8
2022-07	89	79	11,684,594	2	2,069,687	312	226	132,918	2	7,011	5	2,781	148	84	567	6
2022-08	52	43	11,711,797	6	1,056,536	367	242	244,498	6	33,560	10	2,930	132	71	594	14
2022-09	89	78	11,791,742	3	1,876,047	311	195	267,390	9	55,317	4	2,675	136	84	570	16
2022-10	86	76	11,644,094	1	1,625,380	343	223	266,353	9	74,671	4	2,777	135	87	556	21
2022-11	86	78	11,643,674	1	1,810,025	357	211	252,845	5	90,589	2	2,683	149	80	552	9
2022-12	81	72	11,621,509	4	1,704,850	270	179	257,264	3	103,328	5	2,583	125	78	568	7
2023-01	77	69	11,936,594	1	1,763,610	359	181	284,882	6	122,340	3	2,653	138	72	576	16
2023-02	78	72	11,932,500	1	1,831,727	283	195	275,044	4	131,665	3	2,548	121	82	566	9
2023-03	45	42	11,794,287	3	1,587,447	391	176	558,159	7	137,722	4	1,382	119	59	505	16
2023-04	100	91	11,815,297	3	2,442,528	336	282	318,211	7	148,763	5	2,635	143	105	549	12
2023-05	99	92	11,815,500	8	2,511,525	341	308	326,387	4	170,822	11	2,720	142	104	618	6
2023-06	103	94	11,930,797	3	2,585,158	368	329	322,413	5,022	173,141	6	2,946	151	100	612	20
2023-07	111	100	11,929,470	2	2,475,620	423	281	296,993	60	199,534	5	2,816	166	109	604	11
2023-08	108	99	11,701,297	2	2,430,242	457	291	323,026	9	206,814	4	2,840	176	105	619	24
2023-09	120	111	11,938,094	2	2,733,240	444	312	297,203	11	208,460	4	2,794	170	112	643	17
2023-10	121	111	11,848,297	3	2,656,772	417	308	305,045	76	235,442	7	2,767	163	116	617	23
2023-11	131	120	11,936,797	3	2,729,388	433	272	283,485	6	241,965	6	2,806	162	110	655	16
2023-12	119	112	11,899,520	4	3,125,279	511	244	315,932	6	24	6	2,973	189	104	658	16
2024-01	139	130	11,899,204	3	3,308,221	517	285	289,657	10	3,025	4	2,845	193	113	666	29
2024-02	141	131	11,805,924	4	3,482,971	513	321	292,197	280	872	9	2,858	184	122	673	12
2024-03	143	132	11,742,505	5	3,342,025	515	295	292,566	8	543	7	2,792	191	119	687	18
2024-04	145	134	11,670,266	4	3,191,060	492	301	314,184	163	2,951	10	2,838	213	132	691	21
2024-05	149	139	11,692,975	3	3,458,631	467	289	316,181	75	2,401	7	2,774	203	129	688	26
2024-06	148	137	11,740,775	6	3,493,661	697	271	0	49	1,977	12	2,778	221	122	0	24
2024-07	145	136	11,738,177	4	3,563,791	672	257	308,846	359	1,678	9	2,733	199	119	658	24
2024-08	152	140	11,698,051	4	4,037,808	684	294	430,804	27	1,650	7	2,749	211	128	688	25
2024-09	157	147	11,757,767	7	4,006,613	802	270	462,335	296	1,962	16	2,705	215	124	702	80
2024-10	196	182	11,739,170	8	3,628,176	855	355	410,504	167	2,973	16	2,732	232	135	766	82
2024-11	203	185	11,745,777	5	3,730,743	758	230	477,099	175	961	11	2,672	228	105	752	47
2024-12	206	194	11,998,993	4	3,964,443	805	102	490,588	22	5,978	10	2,664	237	67	748	26
Average:	**97.2**	**89.1**	**11,306,358.9**	**2.0**	**2,387,153.0**	**9,543.5**	**160.2**	**189,717.8**	**73.1**	**291,954.6**	**3.8**	**2,573.9**	**127.5**	**73.0**	**465.7**	**10.5**

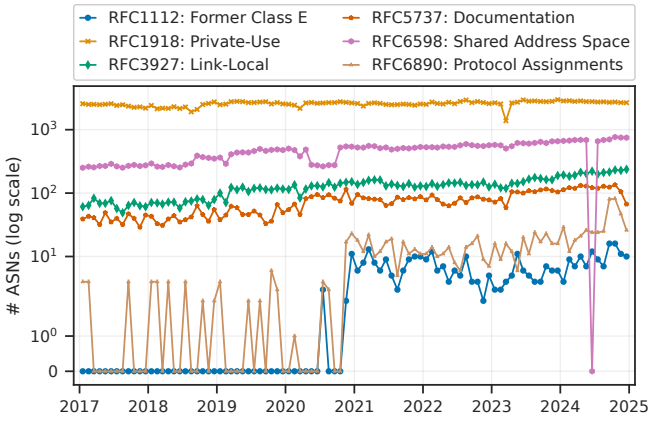

Fig. 5: **BA:** Number of ASNs Transiting Bogons per RFC

TABLE III: **BA:** Breakdown of the Number of Unique ASNs Transiting Bogon IPs per Year and per Bogon Type (RFC)

Year	RFC						Unique per Year
	1112	1918	3927	5737	6598	6890	
2017	0	5,624	180	96	609	8	5,748
2018	0	5,785	225	114	792	12	5,950
2019	0	6,507	306	157	1,078	12	6,754
2020	13	6,489	323	193	1,057	37	6,735
2021	30	6,129	354	196	1,246	60	6,485
2022	23	6,493	333	165	1,338	61	6,865
2023	21	6,514	369	213	1,469	74	6,922
2024	47	6,563	438	254	1,571	169	6,951
Unique per RFC	74	14,896	875	502	3,241	241	15,541

per RFC over the study period (the **BA** case). Results show that most ASNs transiting packets with Bogon addresses do not filter out Private-Use (RFC1918), Shared Address Space (RFC6598), Link-Local (RFC3927) and Documentation (RFC5737) addresses.

For the whole period of the analyzed dataset, we have found 15,541 unique ASes transiting packets with Bogon addresses. Table III reports the breakdown of the number of unique ASNs transiting Bogons per year and per RFC (the **BA** case). Over the 96 months of measurements, we observed a total of 14,896 ASNs transiting Private-Use (RFC1918), 3,241 not filtering Shared Address Space (RFC6598), 875 with Link-Local (RFC3927) addresses and 502 transiting packets coming from Documentation (RFC5737) addresses.

When using the more conservative approach of only recording the AS immediately before the Bogon address (the **BB** case), the number of ASes transiting Bogons is slightly lower, as shown in Table IV, with 14,030 ASNs forwarding packets for Private-Use (RFC1918) addresses, 2,477 forwarding Shared Address Space (RFC6598), 449 allowing packets from Link-Local (RFC3927) addresses, and 107 ASes forwarding packets coming from Documentation (RFC5737) addresses.

Similarly, Table V presents the number of ASes that not only transit packets with Bogon addresses, but can also be confidently identified as using those addresses within their

TABLE IV: **BB:** Breakdown of the Number of Unique ASNs Transiting or Containing Bogon IPs per Year and per Bogon Type (RFC)

Year	RFC						Unique per Year
	1112	1918	3927	5737	6598	6890	
2017	0	5,164	74	23	385	1	5,295
2018	0	5,293	74	15	498	2	5,464
2019	0	6,016	104	23	736	3	6,254
2020	2	6,014	140	31	764	3	6,268
2021	2	5,642	167	33	935	3	6,010
2022	2	6,008	173	33	1,027	5	6,381
2023	1	6,070	179	43	1,121	11	6,497
2024	2	6,091	213	37	1,178	16	6,470
Unique per RFC	3	14,030	449	107	2,477	21	14,672

TABLE V: **BC:** Breakdown of the Number of Unique ASNs Containing Bogon IPs per Year and per Bogon Type (RFC)

Year	RFC						Unique per Year
	1112	1918	3927	5737	6598	6890	
2017	0	2,807	21	4	133	1	2,887
2018	0	2,770	20	5	181	1	2,886
2019	0	3,126	27	5	239	0	3,264
2020	1	3,072	27	4	256	1	3,233
2021	0	2,861	27	6	312	1	3,070
2022	0	3,024	32	6	353	2	3,249
2023	0	3,051	32	7	382	2	3,296
2024	0	3,031	39	9	401	5	3,282
Unique per RFC	1	8,355	122	26	1,012	6	8,805

own infrastructure (the **BC** case). For this case, we identify 8,355 ASes transiting packets with Private-Use (RFC1918) addresses, 1,012 ASes with Shared Address Space (RFC6598), 122 ASes with Link-Local (RFC3927) addresses, and 26 ASes using Documentation (RFC5737) IPs.

The number of unique ASNs transiting packets with Bogon source addresses has steadily increased over the years, possibly due to the increase of Ark VPs or the lack of IPv4 addresses since the IPv4 runout, though the exact number of ASes transiting Bogons varies considerably (see Table II). For instance, the number of ASNs transiting Private-Use (RFC1918) addresses varies from 1,382 to 2,973. Figure 6 shows the empirical Cumulative Distribution Function (eCDF) of ASN occurrence number in the measurements, presented separately for each Bogon type as well as for all types combined. For instance, over the 96 months of measurements, out of the total 14,896 ASNs transiting the Private-Use (RFC1918) addresses, 3,409 (22.89%) ASNs are observed only once. Only 9,894 (66.42%) ASNs are present in more than two measurements, and 1,730 (11.61%) appear in at least half of the measurements. We find 636 (4.27%) ASNs appear in over 86 measurements (>90% of all measurements) and 483 (3.24%) ASNs appear in over 95% of the measurements. The number of ASNs found in all 96 measurements is only 251, representing 1.68% of the total observed ASNs. We also observe significant variations for the other two most popular Bogon address types: Shared Address Space (RFC6598) and Link-Local (RFC3927). These results indicate that while a

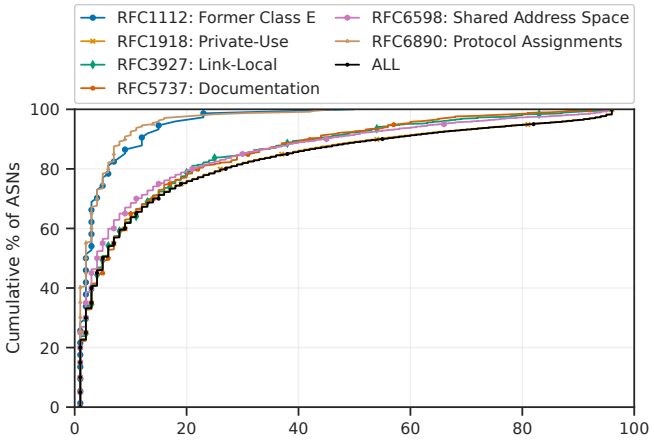

Fig. 6: ASN Occurrences eCDF

large number of unique ASNs are observed, the frequency of appearance varies significantly, with only a portion of ASNs being observed across multiple measurements.

To exemplify this, we built a matrix containing the Jaccard similarity of the sets of ASNs transiting Bogons across different months in 2024 per each type (see Figure 7). As we can see, the Jaccard similarity across two different months for Private-Use (Figure 7a) and Link-Local (Figure 7c) addresses fluctuates around 0.5, slightly decreasing over time, meaning that only around half of ASNs are the same in the two measurements. The Jaccard similarity for Shared Address Space Bogons (Figure 7b) exhibits similar behavior. However, in June 2024, no traceroutes containing this type of Bogon addresses were observed, resulting in the corresponding row and column values being 0.0. For Documentation addresses, we observe a clear decreasing trend over time in the Jaccard similarity, meaning that fewer ASNs remain common (see (Figure 7d)). At the same time, the Jaccard similarity between two different months for Protocol Assignments (Figure 7e) and Former Class E (Figure 7f) addresses is lower than that observed for other bogon address types and fluctuates around 0.2. We attribute this behavior to the smaller number of ASNs that transit these address types.

A possible explanation for this is the limitations of our approach (discussed in Section VI) and the methodology used by CAIDA Ark to collect the traceroutes, which is optimized for finding alternate paths by randomizing destination IPs inside a /24 and performing the traceroute from different vantage points. Indeed, in order for Bogon addresses to be visible as hops, all ASNs in the path to/from the destination need to transit the Bogon packets without filtering them. The use of another vantage point may result in a different path with the ASN performing filtering in it. Thus, the corresponding traceroute will no longer contain Bogon addresses.

These results suggest that ASNs that appear repeatedly in our measurements are more likely to be transiting Bogons and should therefore be scrutinized first. However, the better approach would be to reduce the interval between Ark measurements and focus on ASNs that appear consistently across several most recent measurement rounds.

We also compared the sets of ASNs transiting different types of Bogons across each other. For this task, we chose the *Containment* similarity, an asymmetrical metric that compares the extent to which one set contains elements of the other set. Figure 8a shows the containment similarity matrix of ASNs transiting Bogon packets separated per RFCs for the whole dataset, while Figure 8b reports the numbers only for measurements done in 2024. As expected, the smaller the size of a corresponding set, the higher the probability it will be contained in a set of the larger size. The number of ASNs transiting Bogon packets from the Former Class E (RFC 1112) space is the lowest; therefore, it is not surprising that they exhibit very high containment similarity with the ASNs transiting other types of Bogons. At the same time, the set covering Private-Use (RFC1918) addresses does not fully cover the other datasets. This means that it is worth considering all types of Bogons when evaluating the filtering practices of ASNs.

C. Characterization of ASNs Transiting Bogons

In this section, we characterize the ASNs identified as transiting Bogons (the **BA** case) according to the methodology explained in Section III (see Step ⑤). We concentrate specifically on the results for 2024, as the findings for the preceding periods might be outdated.

Using the ASRank [58] dataset, we associated each identified ASN with its geographical coordinates representing its registration address. Figure 9 presents the scatterplots of ASNs transiting (a) Private-Use (RFC1918) and (b) Shared Address Space (RFC6598) Bogons. The color represents the number of times ASN occurred in our measurements in 2024.

We can draw two conclusions from the figure. First, most ASNs appear in the results occasionally: the majority of dots have a bluish color, representing rare occurrences of the corresponding ASNs across the measurements (see Section IV-B for detailed analysis). As we already noted, such behavior is most probably caused by frequent rotation of the vantage points. Second, most ASes transiting Bogons are located in Europe, Asia and North America. This comes as no surprise because many organizations are also operating in these regions.

To characterize the distribution, we associated ASNs transiting Bogons with the country and Regional Internet Registry (RIR) using the CAIDA AS2ORG [57] dataset. Table VI shows that most of the ASNs identified as transiting Bogons in 2024 (34.73%) are registered with RIPE, while APNIC, ARIN, and LACNIC exhibit very similar numbers.

Table VII lists the Top 10 countries by the number of unique ASNs transiting Bogons in 2024. The USA and Brazil occupy the first two positions with huge margins correspondingly. Interestingly, these findings echo trends observed in the Spoofer dataset, where ASes in the USA and Brazil also rank prominently in terms of the percentage of spoofable ASes [60]. Such alignment suggests a potential correlation with suboptimal filtering practices in these regions.

	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec
Jan	1.00	0.56	0.55	0.54	0.52	0.50	0.50	0.50	0.48	0.47	0.47	0.46
Feb	0.56	1.00	0.55	0.53	0.53	0.52	0.51	0.50	0.49	0.48	0.48	0.47
Mar	0.55	0.55	1.00	0.55	0.55	0.53	0.52	0.50	0.50	0.49	0.48	0.47
Apr	0.54	0.53	0.55	1.00	0.56	0.53	0.52	0.51	0.50	0.51	0.49	0.48
May	0.52	0.53	0.55	0.56	1.00	0.55	0.53	0.52	0.51	0.50	0.50	0.49
Jun	0.50	0.52	0.53	0.53	0.55	1.00	0.56	0.54	0.53	0.52	0.51	0.49
Jul	0.50	0.51	0.52	0.52	0.53	0.56	1.00	0.55	0.53	0.52	0.52	0.50
Aug	0.50	0.50	0.50	0.51	0.52	0.54	0.55	1.00	0.55	0.54	0.53	0.51
Sep	0.48	0.49	0.50	0.50	0.51	0.53	0.53	0.55	1.00	0.55	0.54	0.53
Oct	0.47	0.48	0.49	0.51	0.50	0.52	0.52	0.54	0.55	1.00	0.55	0.54
Nov	0.47	0.48	0.48	0.49	0.50	0.51	0.52	0.53	0.54	0.55	1.00	0.56
Dec	0.46	0.47	0.47	0.48	0.49	0.49	0.50	0.51	0.53	0.54	0.56	1.00

(a) Private-Use (RFC1918)

	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec
Jan	1.00	0.60	0.58	0.55	0.54	0.00	0.48	0.50	0.51	0.46	0.45	0.45
Feb	0.60	1.00	0.61	0.58	0.57	0.00	0.50	0.53	0.53	0.51	0.49	0.47
Mar	0.58	0.61	1.00	0.58	0.58	0.00	0.51	0.53	0.53	0.48	0.48	0.46
Apr	0.55	0.58	0.58	1.00	0.58	0.00	0.50	0.52	0.53	0.50	0.48	0.46
May	0.54	0.57	0.58	0.58	1.00	0.00	0.54	0.55	0.56	0.52	0.51	0.49
Jun	0.00	0.00	0.00	0.00	0.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00
Jul	0.48	0.50	0.51	0.50	0.54	0.00	1.00	0.55	0.55	0.48	0.52	0.48
Aug	0.50	0.53	0.53	0.52	0.55	0.00	0.55	1.00	0.58	0.55	0.55	0.51
Sep	0.51	0.53	0.53	0.53	0.56	0.00	0.55	0.58	1.00	0.57	0.56	0.55
Oct	0.46	0.51	0.48	0.50	0.52	0.00	0.48	0.55	0.57	1.00	0.61	0.60
Nov	0.45	0.49	0.48	0.48	0.51	0.00	0.52	0.55	0.56	0.61	1.00	0.64
Dec	0.45	0.47	0.46	0.46	0.49	0.00	0.48	0.51	0.55	0.60	0.64	1.00

(b) Shared Address Space (RFC6598)

	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec
Jan	1.00	0.63	0.62	0.59	0.55	0.58	0.55	0.48	0.53	0.48	0.49	0.49
Feb	0.63	1.00	0.58	0.54	0.52	0.53	0.54	0.51	0.55	0.49	0.45	0.48
Mar	0.62	0.58	1.00	0.60	0.60	0.63	0.58	0.53	0.55	0.51	0.52	0.50
Apr	0.59	0.54	0.60	1.00	0.62	0.63	0.61	0.56	0.57	0.51	0.49	0.52
May	0.55	0.52	0.60	0.62	1.00	0.65	0.63	0.54	0.65	0.51	0.52	0.55
Jun	0.58	0.53	0.63	0.63	0.65	1.00	0.64	0.63	0.63	0.54	0.55	0.57
Jul	0.55	0.54	0.58	0.61	0.63	0.64	1.00	0.60	0.60	0.52	0.53	0.52
Aug	0.48	0.51	0.53	0.56	0.54	0.63	0.60	1.00	0.62	0.55	0.53	0.53
Sep	0.53	0.55	0.55	0.57	0.65	0.63	0.60	0.62	1.00	0.61	0.61	0.61
Oct	0.48	0.49	0.51	0.51	0.51	0.54	0.52	0.55	0.61	1.00	0.64	0.65
Nov	0.49	0.45	0.52	0.49	0.52	0.55	0.53	0.53	0.61	0.64	1.00	0.67
Dec	0.49	0.48	0.50	0.52	0.55	0.57	0.52	0.53	0.61	0.65	0.67	1.00

(c) Link-Local (RFC3927)

	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec
Jan	1.00	0.81	0.69	0.63	0.62	0.64	0.61	0.51	0.56	0.45	0.38	0.30
Feb	0.81	1.00	0.65	0.64	0.62	0.63	0.60	0.56	0.60	0.47	0.38	0.29
Mar	0.69	0.65	1.00	0.67	0.61	0.63	0.59	0.52	0.51	0.44	0.37	0.26
Apr	0.63	0.64	0.67	1.00	0.68	0.65	0.64	0.62	0.61	0.45	0.42	0.26
May	0.62	0.62	0.61	0.68	1.00	0.66	0.64	0.55	0.58	0.45	0.41	0.27
Jun	0.64	0.63	0.63	0.65	0.66	1.00	0.66	0.55	0.61	0.45	0.45	0.30
Jul	0.61	0.60	0.59	0.64	0.64	0.66	1.00	0.60	0.63	0.50	0.46	0.30
Aug	0.51	0.56	0.52	0.62	0.55	0.55	0.60	1.00	0.67	0.49	0.44	0.28
Sep	0.56	0.60	0.51	0.61	0.58	0.61	0.63	0.67	1.00	0.55	0.50	0.32
Oct	0.45	0.47	0.44	0.45	0.45	0.45	0.50	0.49	0.55	1.00	0.52	0.35
Nov	0.38	0.38	0.37	0.42	0.41	0.45	0.46	0.44	0.50	0.52	1.00	0.39
Dec	0.30	0.29	0.26	0.26	0.27	0.30	0.30	0.28	0.32	0.35	0.39	1.00

(d) Documentation (RFC5737)

	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec
Jan	1.00	0.17	0.21	0.16	0.20	0.13	0.23	0.15	0.18	0.18	0.15	0.22
Feb	0.17	1.00	0.43	0.22	0.27	0.09	0.16	0.12	0.12	0.11	0.13	0.15
Mar	0.21	0.43	1.00	0.22	0.19	0.17	0.24	0.19	0.15	0.16	0.20	0.29
Apr	0.16	0.22	0.22	1.00	0.24	0.18	0.22	0.21	0.16	0.14	0.24	0.21
May	0.20	0.27	0.19	0.24	1.00	0.25	0.16	0.19	0.18	0.16	0.22	0.18
Jun	0.13	0.09	0.17	0.18	0.25	1.00	0.17	0.29	0.14	0.14	0.18	0.19
Jul	0.23	0.16	0.24	0.22	0.16	0.17	1.00	0.20	0.18	0.18	0.18	0.35
Aug	0.15	0.12	0.19	0.21	0.19	0.29	0.20	1.00	0.17	0.23	0.22	0.21
Sep	0.18	0.12	0.15	0.16	0.18	0.14	0.18	0.17	1.00	0.36	0.27	0.14
Oct	0.18	0.11	0.16	0.14	0.16	0.14	0.18	0.23	0.36	1.00	0.26	0.17
Nov	0.15	0.13	0.20	0.24	0.22	0.18	0.18	0.22	0.27	0.26	1.00	0.20
Dec	0.22	0.15	0.29	0.21	0.18	0.19	0.35	0.21	0.14	0.17	0.20	1.00

(e) Protocol Assignments (RFC6890)

	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec
Jan	1.00	0.18	0.38	0.27	0.22	0.33	0.30	0.38	0.25	0.11	0.36	0.08
Feb	0.18	1.00	0.23	0.58	0.07	0.11	0.12	0.14	0.14	0.19	0.18	0.12
Mar	0.38	0.23	1.00	0.13	0.17	0.19	0.33	0.40	0.15	0.10	0.29	0.06
Apr	0.27	0.58	0.13	1.00	0.06	0.16	0.27	0.13	0.18	0.24	0.24	0.05
May	0.22	0.07	0.17	0.06	1.00	0.19	0.45	0.17	0.28	0.21	0.20	0.31
Jun	0.33	0.11	0.19	0.16	0.19	1.00	0.17	0.27	0.27	0.08	0.21	0.10
Jul	0.30	0.12	0.33	0.27	0.45	0.17	1.00	0.33	0.39	0.32	0.33	0.12
Aug	0.38	0.14	0.40	0.13	0.17	0.27	0.33	1.00	0.15	0.10	0.20	0.06
Sep	0.25	0.14	0.15	0.18	0.28	0.27	0.39	0.15	1.00	0.28	0.35	0.13
Oct	0.11	0.19	0.10	0.24	0.21	0.08	0.32	0.10	0.28	1.00	0.35	0.13
Nov	0.36	0.18	0.29	0.24	0.20	0.21	0.33	0.20	0.35	0.35	1.00	0.17
Dec	0.08	0.12	0.06	0.05	0.31	0.10	0.12	0.06	0.13	0.13	0.17	1.00

(f) Former Class E (RFC1112)

Fig. 7: **BA**: Jaccard Similarity of ASNs Transiting Bogon packets of Particular Type across Months in 2024

RFC	1112	1918	3927	5737	6598	6890
1112	1.00	0.99	0.91	0.96	0.97	0.76
1918	0.00	1.00	0.06	0.03	0.18	0.02
3927	0.08	0.95	1.00	0.43	0.76	0.23
5737	0.14	0.97	0.75	1.00	0.92	0.40
6598	0.02	0.82	0.20	0.14	1.00	0.07
6890	0.23	0.98	0.83	0.84	0.94	1.00

(a) Whole Dataset

1112	1918	3927	5737	6598	6890	RFC
1.00	0.98	0.96	0.91	0.98	0.74	1112
0.01	1.00	0.06	0.04	0.19	0.02	1918
0.10	0.94	1.00	0.45	0.78	0.31	3927
0.17	0.94	0.78	1.00	0.93	0.50	5737
0.03	0.77	0.22	0.15	1.00	0.10	6598
0.21	0.95	0.79	0.75	0.90	1.00	6890

(b) In 2024

Fig. 8: **BA:** Containment Similarity of ASN Sets Transiting Bogons per RFC

Table VII also provides a breakdown of the number of unique ASNs by bogon type, while Figure 10 presents the data for the entire world as colormaps.

The majority of ASNs transiting Bogons (76.94%) are in the Computer and Information Technology category, with 65.49% being Internet Service Providers (see Table VIII). This is not an unusual finding, as it is expected that ISPs provide transit services, similar to NRENs (National Research and Education Networks) providing transit to Educational institutions.

TABLE VI: **BA:** RIR for ASNs Transiting Bogons in 2024

RIR	ASNs	%
RIPE	2,414	34.73%
APNIC	1,431	20.59%
ARIN	1,392	20.03%
LACNIC	1,334	19.19%
AFRINIC	364	5.24%
Not found	16	0.23%
Total ASNs	**6,951**	**100.00%**

V. Validation

A. *Cross-check with the Spoofer Dataset*

The Spoofer dataset contains results of 988,254 individual measurements performed during our analysis, related to 9,051

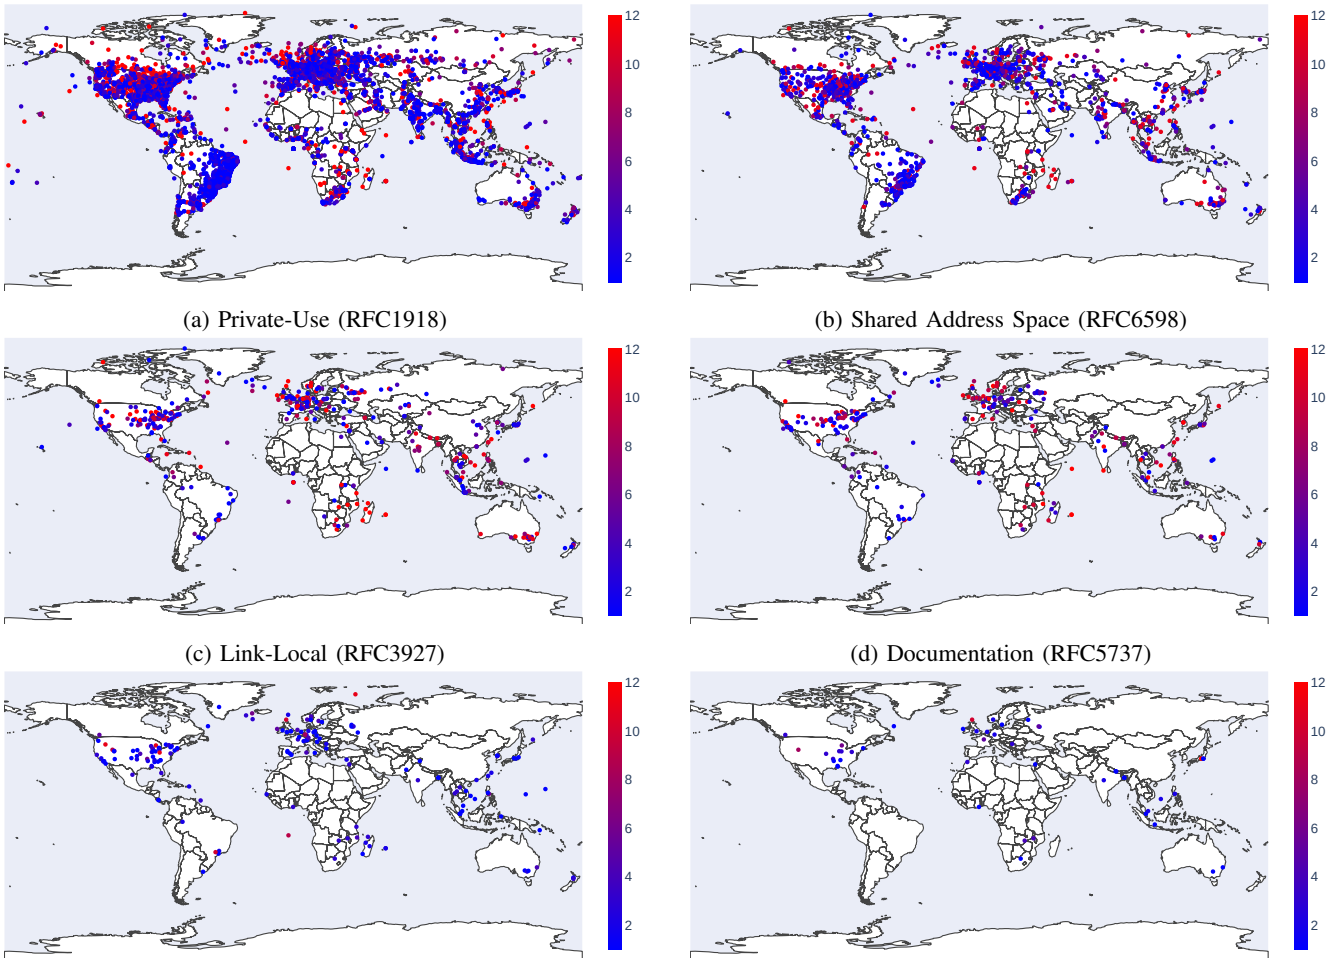

(a) Private-Use (RFC1918)

(b) Shared Address Space (RFC6598)

(c) Link-Local (RFC3927)

(d) Documentation (RFC5737)

(e) Protocol Assignments (RFC6890)

(f) Former Class E (RFC1112)

Fig. 9: **BA:** Map of ASNs Occurrences Transiting Bogons of Particular Type in 2024

TABLE VII: **BA:** Top 10 Countries by Number of ASNs Transiting Bogons

Country Code	# Unique ASNs	# Unique ASNs per RFCs					
		1112	1918	3927	5737	6598	6890
US	**1207**	14	1122	120	83	335	53
BR	**938**	0	869	11	8	197	3
RU	**437**	3	421	27	10	65	5
ID	**288**	1	279	8	1	32	0
GB	**181**	1	160	18	9	55	7
PL	**175**	0	171	5	3	20	0
CA	**173**	0	167	3	2	34	2
DE	**171**	1	158	13	3	40	7
BD	**170**	2	170	2	2	12	2
IT	**155**	0	150	3	2	25	2

TABLE VIII: **BA:** Categories of ASNs Identified as Transiting Bogons in 2024 according to ASdb

Category	**ASNs**	**%**
Computer and Information Technology - Internet Service Provider (ISP)	4,552	65.49%
Computer and Information Technology - (no second category found)	488	7.02%
Computer and Information Technology - Hosting and Cloud Provider	308	4.43%
Education and Research - Colleges, Universities, and Professional Schools	175	2.52%
Other	1,232	17.72%
Not found	196	2.82%
Total ASNs 2024	**6,951**	**100%**

unique ASNs. This number corresponds to roughly 11% of all active ASNs, indicating low coverage for the dataset. Moreover, it is considerably lower than the number of unique ASNs identified in this work as transiting Bogons (15,541).

While handling IP spoofing and filtering Bogon addresses entails similar but slightly different technical approaches, the correlation of these network hygiene practices provides interesting results. Out of 9,051 unique ASNs tested by Spoofer, 5,707 are non-spoofable (the corresponding routable spoofed packets are `blocked`), and 2,880 are spoofable (the matching `routedspoof` packets are `received`). Out of the spoofable and not spoofable ASNs, 1,898 have both spoofable and not spoofable results. For the rest 464 ASNs in the Spoofer dataset, the respective `routedspoof` packets

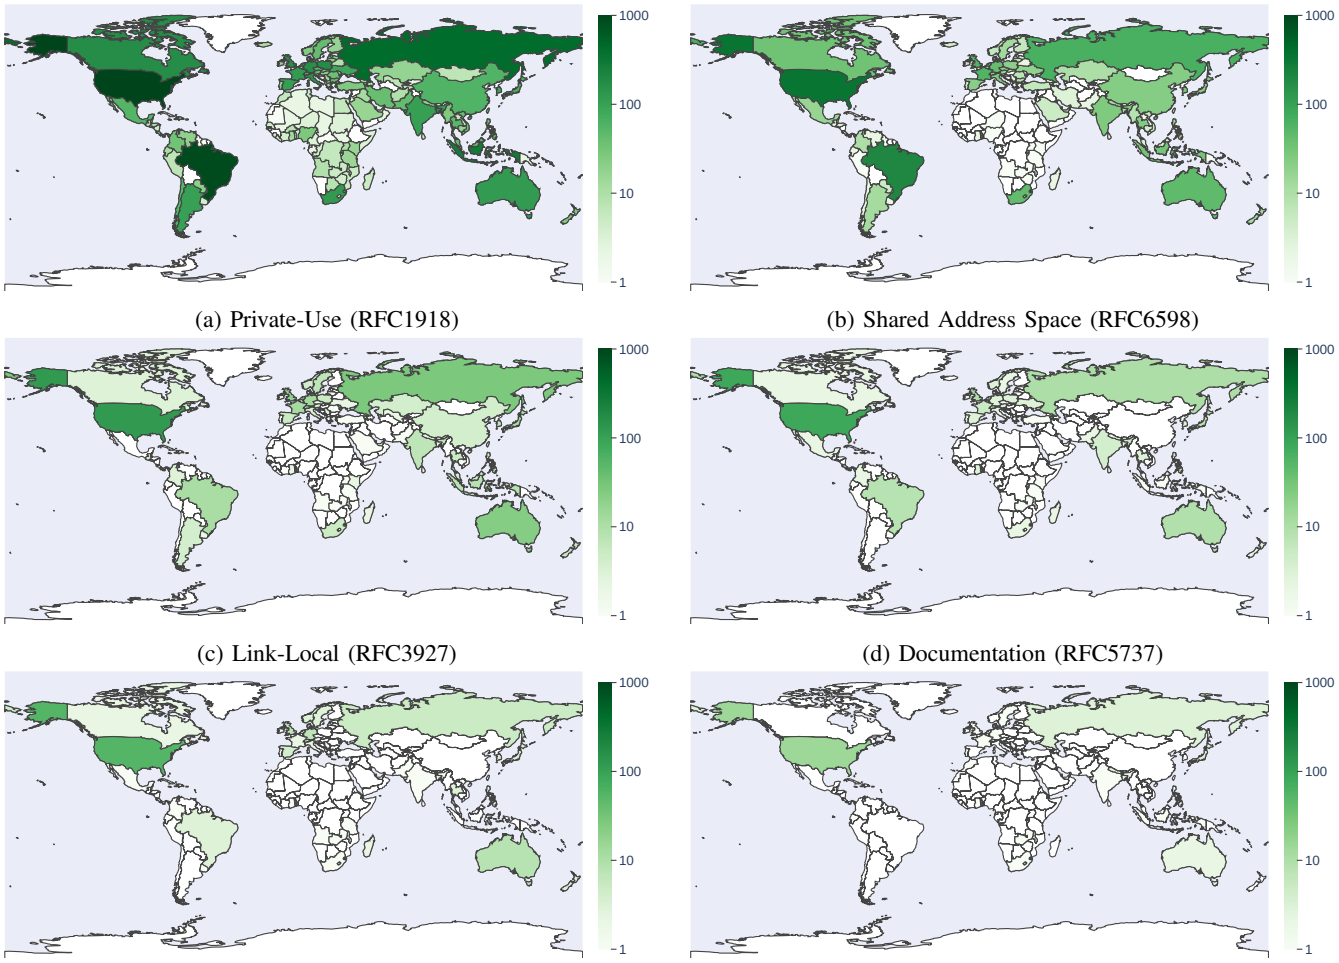

(a) Private-Use (RFC1918)

(b) Shared Address Space (RFC6598)

(c) Link-Local (RFC3927)

(d) Documentation (RFC5737)

(e) Protocol Assignments (RFC6890)

(f) Former Class E (RFC1112)

Fig. 10: **BA:** Colormap of the Number of ASNs per Country Transiting Bogons of Particular Type in 2024

are either `rewritten`, `unknown` or `N/A`, with the latter signifying that the measurement is for IPv6 only.

As we describe in Section III, for each identified ASN transiting Bogon packets, we look for the data about it in the Spoofer dataset during a period of six months before the corresponding measurement. We report the ASN as *spoofable* if the data is found and the corresponding entry for `routedspoof` from this ASN is `received`. Respectively, we consider the ASN as *non-spoofable* if the data is found and the matching `routedspoof` packet to this ASN is `blocked`.

Table IX reports the results. In the "Identified Unique ASNs", we report the number of unique ASNs identified as transiting Bogon packets of the corresponding "Bogon Type". The "# in Spoofer" shows how many of those ASNs are also found in the Spoofer dataset. The next three columns, namely "Only Spoofable", "Only Non-Spoofable" and "Both Spoofable Non-Spoofable", report the numbers and percentages of the ASNs, all measurements for which are either *spoofable*, *non-spoofable* or *both* correspondingly.

For instance, the analysis reveals that out of 14,896 ASNs identified as transiting Private-Use (RFC1918) addresses, 2,529 are also found in the Spoofer dataset. Of them, 358 (14.16%) are spoofable in all the corresponding measurements. This means that, with a high confidence level, we can affirm the corresponding ASNs as not implementing current best practices for SAV. Interestingly, 1,333 (52.71%) ASNs are found only non-spoofable, while they transit Bogon packets, and 725 (28.67%) ASNs are found as both spoofable and non-spoofable in different measurements. There could be multiple reasons why the corresponding ASNs are identified as non-spoofable. First, this may show that the SAV best practices are implemented only in some parts of that network. Second, the issue might have been fixed sometime after it was detected, thus not appearing in the intersection. Moreover, such a result might also be caused by issues in the Spoofer dataset that we cannot identify.

B. Cross-check with the MANRS Dataset

We also cross-check our findings with the MANRS for Network Operators dataset [52], which contains 1,072 entries. However, the "ASNs" column may list multiple ASNs in one entry, thus, describing in total 1,419 unique ASNs.

TABLE IX: **BA:** Comparison of ASNs Identified as Transiting Bogons with the CAIDA Spoofer Data (S - Spoofable, NS - Non-Spoofable)

Bogon Type	# ASNs	# in Spoofer	Only S		Only NS		Both S & NS	
			#	%	#	%	#	%
RFC1112	74	30	1	3.33%	20	66.67%	8	26.67%
RFC1918	14,896	2,529	358	14.16%	1,333	52.71%	725	28.67%
RFC3927	875	309	25	8.09%	159	51.46%	116	37.54%
RFC5737	502	226	13	5.75%	125	55.31%	85	37.61%
RFC6598	3,241	811	97	11.96%	408	50.31%	283	34.90%
RFC6890	241	88	5	5.68%	56	63.64%	26	29.55%

TABLE X: **BA:** Comparison of ASNs Identified as Transiting Bogons in 2024 with the MANRS Data (Conformance: C - conformant, NC - non-conformant)

Members	Conf.	# Unique ASNs	# Unique ASNs per RFC					
			1112	**1918**	**3927**	**5737**	**6598**	**6890**
All	**C**	258	12	244	64	58	128	28
	NC	154	5	142	23	14	64	10
Before 2024	**C**	231	11	217	60	55	117	27
	NC	129	5	118	20	12	53	9

As we explained in Section III (see Step ⑦), although the MANRS `anti_spoofing` data is a derivative from the Spoofer dataset, there is one core difference between them, namely the commitment of the MANRS members to filter out spoofed packets. So, as the MANRS dataset contains the snapshot of the latest measurement results (in our case, for July 2025), we compared them with the list of ASNs transiting Bogon packets in 2024. Table X reports on the comparison results. As we can see, among 6,951 unique ASNs identified as transiting Bogons in 2024, the overlap with MANRS participating networks is only 412. The majority of these ASNs, 258, claim to be conformant with the MANRS `anti_spoofing` action. It is possible that due to the difference in the datasets' collection time, these ASNs were non-conformant in 2024 but have fixed their networks by the last MANRS measurement. In the future, we plan to collect the data regularly and compare the results collected at a closer timeframe.

To draw conclusions, we separately report on the results for the participants approved as MANRS members before January 1, 2024. We assume that by the start of 2024, these participants would have been notified and implemented all best practices, including Bogon filtering. However, 231 ASNs marked as conforming to the MANRS `anti_spoofing` action were still transiting packets with Bogon addresses in 2024.

VI. Measurement Considerations

Several important measurement considerations must be taken into account when interpreting the results of our study. These include potential biases introduced by the nature of used data, limitations of vantage point coverage, and the temporal scope of observations, all of which may influence the representativeness and accuracy of the findings.

Nature of Used Data. Utilizing traceroute data presents common shortcomings. First, some routers may be configured to avoid responding to traceroute probes or rate limit the responses, while others may respond not from the address of the receiving interface. Second, our findings are based on an incomplete view of the Internet topology. Internet-wide traceroute measurements, such as those performed by the CAIDA Ark project, are limited in their ability to uncover network paths. It is not feasible to trace all IP addresses from each vantage point during every measurement cycle. These limitations make it hard to reproduce or pinpoint when specific changes happened [47].

Moreover, traceroute probes can only observe the forward data plane path from the source vantage point to the destination, while the responses may take a different return path. This asymmetry could be caused by intentional traffic engineering or unintentional routing decisions by the BGP Best Path selection algorithm. Many networks deploy multiple equal-cost paths to destinations, with individual packets load-balanced across alternatives. On any given traceroute probe, only one of potentially many paths is visible; different probes to the same destination may observe entirely different AS sequences. An AS might filter Bogon packets on the inbound path while failing to filter them on the outbound path, or vice versa, potentially failing to identify ASes that are not properly filtering Bogons due to a lack of visibility for the reverse path.

Several studies have examined asymmetry in Internet routing. In 2005, He et al. [61] showed that at least 14% of route pairs exhibited AS-level asymmetry, though this percentage varied considerably depending on the employed measurement infrastructures (or datasets). De Vries et al. [62] found that in 2015 only 12.6% of path tuples were identical at the hop level, though, the relative edit distance between paths remained low (about 6% when measured within 24 hours). Moreover, the author acknowledged that the AS-level asymmetry was considerably lower. More recently in 2022, Vermeulen et al. [63] reported that 47% of paths compared at the AS level are asymmetric. These studies indicate that route asymmetry may indeed influence our findings. However, to mitigate the influence of this factor, we can focus on only those ASes that appear consistently across measurements (see Section IV-B). Moreover, we can prioritize our BC ("sandwich") category, where the confidence is higher because the same AS appears both before and after a Bogon address.

Vantage Point Bias. Our use of CAIDA Ark for traceroute measurements, enriched with BGP data from RIPE RIS and RouteViews, provides global coverage but is subject to bias due to measurement infrastructure limitations. While route collectors are strategically distributed across Europe, North America, Asia, South America, and Africa, yet the geographic concentration remains skewed toward developed regions with high AS density and mature Internet infrastructure. Moreover, route collectors are predominantly deployed at Internet eXchange Points (IXPs), creating bias toward well-connected ASes (particularly tier-1 providers and content delivery networks) while underrepresenting regional operators and small autonomous systems. A comprehensive analysis of Internet

measurement infrastructure demonstrates that both RIPE RIS and RouteViews exhibit significant topological and geographic skew [64]. CAIDA Ark's vantage points are similarly distributed, with concentration in regions offering reliable network operators willing to host monitoring infrastructure. As presented in Section IV-C, 34.73% of ASes transiting Bogons in 2024 are registered with RIPE (Europe), compared to only 5.24% registered with AFRINIC (Africa). This imbalance may indicate either a higher rate of non-compliance with filtering best practices among European ISPs, or, more likely, reduced measurement visibility in developing regions dominated by smaller ISPs. This geographic and topological bias implies that our findings robustly characterize filtering practices among well-connected, primarily ISP-tier networks in developed regions, while providing limited insight into regional operators or networks in developing economies.

Temporal Bias. Our measurements are also subject to temporal bias. Consider an AS implementing Bogon filtering on 9 of 10 Equal-Cost Multipath (ECMP) paths but neglecting the 10th: our measurement captures this failure with only 10% probability per probe. Over multiple measurement rounds, we expect eventual detection, but transient routing misconfigurations lasting hours or days between our monthly measurement cycles may remain invisible. The Jaccard similarity analysis presented in Section IV-B revealed that the overlap of ASes transiting Bogons between consecutive months in 2024 fluctuated around 0.50 for RFC1918 addresses and varied substantially for other Bogon types, consistent with load-balanced path diversity rather than systematic filtering policies [65].

VII. Discussion

A. Key Considerations

Our analysis revealed the lack of filtering packets from Bogon addresses in more than $15{,}500$ ASes, creating an impression of widespread violations. Some of the reasons for such high numbers are due to the biases discussed in Section VI. Still, the presented results are based on actual observations of the networks and thus reflect a sample of real-world scenarios. However, we try to mitigate the biases through multiple strategies: (1) aggregating observations across 96 monthly measurement cycles to identify persistent rather than transient violations; (2) applying a conservative persistency threshold, flagging ASes as violators only when Bogon transit appears in $\geq 90\%$ of measurements; and (3) separately analyzing the BA (all ASes on the path before Bogon), BB (AS immediately before Bogon), and BC (AS appearing before and after Bogon) cases to distinguish systematic use from transient artifacts.

However, the presented considerations show the need to develop different approaches and tools beyond traceroute measurements for identifying networks not filtering Bogon addresses. Incorporating multiple data sources and methodologies, such as deterministic active probing techniques, could improve the detection of ASNs transiting packets from Bogon addresses. Moreover, collaboration between network operators and researchers is crucial to address the limitations of traceroute measurements and improve the accuracy of Internet topology mapping efforts.

As networks grow and the transition to IPv6 does not happen at the same pace, the need to assign IPv4 addresses to equipment increases. With a lack of IPv4 addresses, networks will use more Private-Use (RFC1918) or Shared Address Space (RFC6598) addresses for their infrastructures. The risks of accepting and forwarding these packets without proper filtering also increase. However, as for spoofed packets, the benefits of filtering them are mainly for the neighboring ASes, while the difficulties and increased costs generated by the implementation are sustained by the ASes not directly benefiting. Thus, individual ASes might not see reasons for implementing these measures.

To support and standardize efforts in routing security, initiatives such as MANRS, encourage network operators and equipment vendors to adopt best practices, including Source Address Validation (SAV). However, despite the interest demonstrated by MANRS members through their participation in the initiative, the presence of Bogon addresses visible in traceroutes crossing AS borders shows that some of these networks lack basic security measures. This underscores the substantial work needed to improve routing security and implement effective anti-spoofing measures across all networks on the Internet.

B. Practical Mitigation Recommendations

Adopt Source Address Validation (BCP 38 / BCP 84). ISPs should implement ingress and egress filtering as defined in BCP 38 and BCP 84. Bogon prefixes do not have valid routes outside the AS, and Private-Use (RFC1918) address space is not globally unique; therefore, any packets crossing the AS border using these addresses as sources should be dropped. The MANRS initiative provides practical guidance for implementing these controls, and ISPs can verify their filtering using tools such as the CAIDA Spoofer project or the MANRS Observatory.

Continuous Traffic Monitoring. Flow-based monitoring at network edges can detect packets with Bogon source addresses in near real time. This allows operators to identify the originating AS and coordinate directly with the responsible network to resolve the issue.

Operator Notification Mechanisms. Establishing automated systems to notify network operators when Bogon-sourced traffic is observed transiting their AS can help shorten remediation timelines. Even if previous studies suggest limited effectiveness [66], [67], timely visibility remains valuable for operators who intend to maintain proper filtering.

Public Reporting and Transparency. Public disclosure of persistent filtering violations (sometimes referred to as "naming and shaming") has demonstrated limited but measurable positive impact [67]. However, disclosure alone is not sufficient; it should be paired with incentives, community engagement, and operational support for remediation.

Voluntary, Industry-Led Adoption of Best Practices. To avoid inconsistent or overly restrictive regulatory intervention across regions, ISPs should collaboratively and voluntarily adopt network hygiene best practices. Broad industry alignment reduces the risk of fragmented policies and ensures stable global interoperability.

VIII. Related work

Bogon addresses often appear in Internet measurements. Despite their commonality, academic papers tend to perceive them as anomalies or measurement errors [68]. Also, attributing them to specific network infrastructures without additional data is a difficult task.

Bogon Addresses. Much of the literature on Bogon addresses is dedicated to outlining strategies for their filtering. For example, filtering routes based on assignments from the Regional Internet Registries in order to reduce the routing table size [69]. A study on Bogon routes [70] looks for such routes in the BGP Global Routing Table. In [26], the usage of former Class E space, 240.0.0.0/4, is identified inside the networks of large companies such as Amazon and Verizon Business. Our research also looks for former Class E address space, but as seen from vantage points outside the AS using them.

Topology Discovery with Traceroutes. Traceroutes are frequently used in Internet measurements in order to uncover network paths otherwise not visible, and large datasets of such traceroutes exist (Ark [47], Atlas [71]). The CAIDA Ark dataset is found to achieve the most comprehensive coverage of node discovery by [72]. One type of information that can be found using traceroutes is routing loops. [15] uses routing loops to infer the ability to send spoofed packets from inside one network based on the loops found between interconnecting ASes. Loops as a security risk for amplification and other types of attacks are found by [68], while their research excludes Private-Use (RFC1918) addresses on purpose. In [73], traceroute information is used to identify networks squatting IP addresses assigned to other organizations; this work is similar to the one proposed in our paper; however, because the authors are looking for IP addresses that are assigned to organizations, no implications about SAV can be inferred, as the squatted IPs are expected to be found routed on the Internet.

Source Address Validation. As SAV is a problem affecting the correct operation of the Internet, in current research, various ways of detecting networks not implementing SAV are developed, some are based on active measurements [15], using spoofed packets [16] with various protocols/IP header options [74], some more intrusive using BGP poisoning in combination with passive honeypots [75]. The CAIDA Spoofer project offers active measurements run by volunteers to identify spoofable networks [9]; in our research, we also cross-check our findings using this dataset. A study [40] of flow data from a large European IXP finds 72% of members send packets with Bogon sources, concluding that the majority of members do not filter outbound traffic.

IX. Conclusions

Bogon addresses frequently appear in Internet measurements but are often dismissed as errors or anomalies. Our paper shows they are not anomalies but common results of poor network hygiene that researchers should account for and operators should fix. Our paper shows the extent of this problem, which has been previously noticed.

Across the 96 measurements collected over eight years, we observe a clear upward trend in ASes transiting packets with Bogon source addresses, particularly Private-Use (RFC1918), Link-Local (RFC3927), and Shared Address Space (RFC6598). This rise is likely driven by growing reliance on private addressing due to IPv4 exhaustion, combined with insufficient filtering at AS interconnection points.

Using available datasets to infer spoofing capability of ASes that transit Bogon traffic, we find fewer spoofable ASes than non-spoofable ones. Some ASes appear in both categories, indicating that SAV policies may vary by prefix or network segment. However, the data is limited by the Spoofer project's coverage, and Bogon transiting may stem from misconfigurations unrelated to SAV. Given the small number of tested ASes, we cannot determine whether SAV is applied to non-private address space within ASes that transit Bogons.

References

[1] G. Jones (Ed.), "Operational Security Requirements for Large Internet Service Provider (ISP) IP Network Infrastructure," RFC 3871 (Informational), RFC Editor, Fremont, CA, USA, Sep. 2004, updated by RFC 8996. [Online]. Available: https://www.rfc-editor.org/rfc/rfc3871.txt

[2] V. Paxson, "An Analysis of Using Reflectors for Distributed Denial-of-Service Attacks," *ACM SIGCOMM Computer Communication Review*, vol. 31, no. 3, pp. 38–47, jul 2001.

[3] C. Rossow, "Amplification Hell: Revisiting Network Protocols for DDoS Abuse," in *Network and Distributed System Security Symposium*, 2014.

[4] R. Anghel, S. Vetrivel, E. Turcios Rodriguez, K. Sameshima, D. Makita, K. Yoshioka, C. H. Gañán, and Y. Zhauniarovich, "Peering into the darkness: The use of utrs in combating ddos attacks," in *European Symposium on Research in Computer Security*, 9 2023, pp. 23–41.

[5] H. Wang, C. Jin, and K. G. Shin, "Defense against spoofed IP traffic using hop-count filtering," *IEEE/ACM Trans. Netw.*, vol. 15, no. 1, pp. 40–53, feb 2007.

[6] A. Yaar, A. Perrig, and D. Song, "StackPi: New packet marking and filtering mechanisms for DDoS and IP spoofing defense," *IEEE Journal on Selected Areas in Communications*, vol. 24, no. 10, pp. 1853–1863, 2006.

[7] S. T. Zargar, J. Joshi, and D. Tipper, "A survey of defense mechanisms against distributed denial of service (DDoS) flooding attacks," *IEEE Communications Surveys & Tutorials*, vol. 15, no. 4, pp. 2046–2069, 2013.

[8] M. Luckie, R. Beverly, R. Koga, K. Keys, J. A. Kroll, and K. Claffy, "Network Hygiene, Incentives, and Regulation: Deployment of Source Address Validation in the Internet," in *ACM SIGSAC Conference on Computer and Communications Security*, 2019, pp. 465–480.

[9] R. Beverly, A. Berger, Y. Hyun, and K. Claffy, "Understanding the Efficacy of Deployed Internet Source Address Validation Filtering," in *ACM SIGCOMM Internet Measurement Conference*, 2009, pp. 356–369.

[10] Cybersecurity and Infrastructure Security Agency, "Principles and Approaches for Security-by- Design and -Default," Cybersecurity and Infrastructure Security Agency, 2023. [Online]. Available: https://www.cisa.gov/sites/default/files/2023-04/principles_approaches_for_security-by-design-default_508_0.pdf

[11] P. Ferguson and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing," RFC 2827 (Best Current Practice), RFC Editor, Fremont, CA, USA, May 2000, updated by RFC 3704. [Online]. Available: https://www.rfc-editor.org/rfc/rfc2827.txt

[12] F. Baker and P. Savola, "Ingress Filtering for Multihomed Networks," RFC 3704 (Best Current Practice), RFC Editor, Fremont, CA, USA, Mar. 2004, updated by RFC 8704. [Online]. Available: https://www.rfc-editor.org/rfc/rfc3704.txt

[13] K. Sriram, D. Montgomery, and J. Haas, "Enhanced Feasible-Path Unicast Reverse Path Forwarding," RFC 8704 (Best Current Practice), RFC Editor, Fremont, CA, USA, Feb. 2020. [Online]. Available: https://www.rfc-editor.org/rfc/rfc8704.txt

[14] R. Beverly, M. Luckie, R. Koga, and K. Keys, "Spoofer API," https://catalog.caida.org/dataset/spoofer_public, dates used: June 2016 - December 2024.

[15] Q. Lone, M. Luckie, M. Korczyński, and M. Van Eeten, "Using loops observed in traceroute to infer the ability to spoof," in *International Conference on Passive and Active Measurement*, 2017, pp. 229–241.

[16] M. Korczyński, Y. Nosyk, Q. Lone, M. Skwarek, B. Jonglez, and A. Duda, "Don't forget to lock the front door! inferring the deployment of source address validation of inbound traffic," in *International Conference on Passive and Active Measurement*, 2020, pp. 107–121.

[17] Internet Society, "MANRS programs," Internet Society, 2021. [Online]. Available: https://manrs.org/programs/

[18] Q. Deng, J. Pu, Z. Tan, Z. Qian, and S. V. Krishnamurthy, "Beyond the Horizon: Uncovering Hosts and Services Behind Misconfigured Firewalls," in *IEEE Symposium on Security and Privacy (SP)*. IEEE, 2025, pp. 1770–1788.

[19] D. Tatang, C. Schneider, and T. Holz, "Large-Scale Analysis of Infrastructure-Leaking DNS Servers," in *Detection of Intrusions and Malware, and Vulnerability Assessment*. Cham: Springer International Publishing, 2019, pp. 353–373.

[20] A. Bechtsoudis and N. Sklavos, "Aiming at Higher Network Security through Extensive Penetration Tests," *IEEE Latin America Transactions*, vol. 10, no. 3, pp. 1752–1756, 2012.

[21] Y. Rekhter, B. Moskowitz, D. Karrenberg, G. J. de Groot, and E. Lear, "Address Allocation for Private Internets," RFC 1918 (Best Current Practice), RFC Editor, Fremont, CA, USA, Feb. 1996, updated by RFC 6761. [Online]. Available: https://www.rfc-editor.org/rfc/rfc1918.txt

[22] J. Weil, V. Kuarsingh, C. Donley, C. Liljenstolpe, and M. Azinger, "IANA-Reserved IPv4 Prefix for Shared Address Space," RFC 6598 (Best Current Practice), RFC Editor, Fremont, CA, USA, Apr. 2012. [Online]. Available: https://www.rfc-editor.org/rfc/rfc6598.txt

[23] S. Cheshire, B. Aboba, and E. Guttman, "Dynamic Configuration of IPv4 Link-Local Addresses," RFC 3927 (Proposed Standard), RFC Editor, Fremont, CA, USA, May 2005. [Online]. Available: https://www.rfc-editor.org/rfc/rfc3927.txt

[24] IANA, "IANA IPv4 Special-Purpose Address Registry," 2024. [Online]. Available: https://www.iana.org/assignments/iana-ipv4-special-registry/iana-ipv4-special-registry.xhtml

[25] S. Deering, "Host extensions for IP multicasting," RFC 1112 (Internet Standard), RFC Editor, Fremont, CA, USA, Aug. 1989, updated by RFC 2236. [Online]. Available: https://www.rfc-editor.org/rfc/rfc1112.txt

[26] Q. Lone, "240/4 as seen by RIPE Atlas," August 2022. [Online]. Available: https://labs.ripe.net/author/qasim-lone/2404-as-seen-by-ripe-atlas/

[27] R. Braden (Ed.), "Requirements for Internet Hosts - Communication Layers," RFC 1122 (Internet Standard), RFC Editor, Fremont, CA, USA, Oct. 1989, updated by RFCs 1349, 4379, 5884, 6093, 6298, 6633, 6864, 8029, 9293. [Online]. Available: https://www.rfc-editor.org/rfc/rfc1122.txt

[28] J. Arkko, M. Cotton, and L. Vegoda, "IPv4 Address Blocks Reserved for Documentation," RFC 5737 (Informational), RFC Editor, Fremont, CA, USA, Jan. 2010. [Online]. Available: https://www.rfc-editor.org/rfc/rfc5737.txt

[29] M. Cotton, L. Vegoda, R. Bonica (Ed.), and B. Haberman, "Special-Purpose IP Address Registries," RFC 6890 (Best Current Practice), RFC Editor, Fremont, CA, USA, Apr. 2013, updated by RFC 8190. [Online]. Available: https://www.rfc-editor.org/rfc/rfc6890.txt

[30] S. Cheshire and D. Schinazi, "Special Use Domain Name 'ipv4only.arpa'," RFC 8880 (Proposed Standard), RFC Editor, Fremont, CA, USA, Aug. 2020. [Online]. Available: https://www.rfc-editor.org/rfc/rfc8880.txt

[31] T. Savolainen, J. Korhonen, and D. Wing, "Discovery of the IPv6 Prefix Used for IPv6 Address Synthesis," RFC 7050 (Proposed Standard), RFC Editor, Fremont, CA, USA, Nov. 2013, updated by RFC 8880. [Online]. Available: https://www.rfc-editor.org/rfc/rfc7050.txt

[32] O. Troan and B. Carpenter (Ed.), "Deprecating the Anycast Prefix for 6to4 Relay Routers," RFC 7526 (Best Current Practice), RFC Editor, Fremont, CA, USA, May 2015. [Online]. Available: https://www.rfc-editor.org/rfc/rfc7526.txt

[33] C. Systems, "Cisco IOS and IOS XE Software SNMP Denial of Service Vulnerability," September 2025. [Online]. Available: https://sec.cloudapps.cisco.com/security/center/content/CiscoSecurityAdvisory/cisco-sa-snmp-x4LPhte

[34] T. M. Research, "Operation Zero Disco: Attackers Exploit Cisco SNMP Vulnerability to Deploy Rootkits," October 2025. [Online]. Available: https://www.trendmicro.com/en_us/research/25/j/operation-zero-disco-cisco-snmp-vulnerability-exploit.html

[35] Censys, "CVE-2017-6742 Actively Exploited SNMP Vulnerability on Cisco Routers," Censys, 2023. [Online]. Available: https://censys.com/blog/cve-2017-6742-actively-exploited-snmp-vulnerability-on-cisco-routers

[36] Dave Phelan, "RIPE91: Filter Your Things or No I don't want to use your NTP Server," 2025. [Online]. Available: https://pretalx.ripe.net/media/ripe91/submissions/BXAUMQ/resources/Filter_your_things-Da_nAvfmSx.pdf

[37] J. Graham-Cumming, "Understanding and mitigating NTP-based DDoS attacks," January 2014. [Online]. Available: https://blog.cloudflare.com/understanding-and-mitigating-ntp-based-ddos-attacks/

[38] J. J. Gondim, R. de Oliveira Albuquerque, and A. L. Sandoval Orozco, "Mirror saturation in amplified reflection Distributed Denial of Service: A case of study using SNMP, SSDP, NTP and DNS protocols," *Future Generation Computer Systems*, vol. 108, pp. 68–81, 2020.

[39] MITRE Corporation, "System Network Configuration Discovery, Technique T1016," October 2025. [Online]. Available: https://attack.mitre.org/techniques/T1016/

[40] F. Lichtblau, F. Streibelt, T. Krüger, P. Richter, and A. Feldmann, "Detection, classification, and analysis of inter-domain traffic with spoofed source IP addresses," in *ACM Internet Measurement Conference*, 2017, pp. 86–99.

[41] Norton, "IP spoofing: What is it and how does it work?" March 2025. [Online]. Available: https://us.norton.com/blog/malware/what-is-ip-spoofing

[42] Y. Zhauniarovich and P. Dodia, "Sorting the Garbage: Filtering Out DRDoS Amplification Traffic in ISP Networks," in *IEEE Conference on Network Softwarization (NetSoft)*, 2019, pp. 142–150.

[43] Team Cymru, "Decoding Bogons: A Senior Stakeholder's Expert Insight," March 2025. [Online]. Available: https://www.team-cymru.com/post/unravelling-the-mystery-of-bogons-a-senior-stakeholder-and-it-professional-guide

[44] Cisco Systems, Inc., *Security Events and Alarm Categories - Version 7.4.2 (DV 2.1)*, Cisco Systems, Inc., 2023. [Online]. Available: https://www.cisco.com/c/dam/en/us/td/docs/security/stealthwatch/management_console/securit_events_alarm_categories/7_4_2_Security_Events_and_Alarm_Categories_DV_2_1.pdf

[45] A. Kirkham, "Issues with Private IP Addressing in the Internet," Internet Engineering Task Force (IETF) — Network Working Group, Internet-Draft draft-kirkham-private-ip-sp-cores-08, Nov. 2011. [Online]. Available: https://www.ietf.org/archive/id/draft-kirkham-private-ip-sp-cores-08.html

[46] CAIDA, "The IPv4 routed /24 topology dataset," CAIDA, dates used: January 2017 - December 2024. [Online]. Available: https://www.caida.org/catalog/datasets/ipv4_routed_24_topology_dataset/

[47] Y. Hyun, "Archipelago (Ark) Measurement Infrastructure," in *CAIDA-WIDE Workshop*, 2006. [Online]. Available: https://www.caida.org/projects/ark/

[48] M. Luckie, "Scamper: A scalable and extensible packet prober for active measurement of the internet," in *ACM SIGCOMM Internet Measurement Conference*, 2010, pp. 239–245.

[49] RIPE NCC, "Routing information service (RIS)." [Online]. Available: https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris/

[50] University of Oregon, "Route views project." [Online]. Available: http://www.routeviews.org/

[51] L. Blunk, M. Karir, and C. Labovitz, "Multi-Threaded Routing Toolkit (MRT) Routing Information Export Format," RFC 6396 (Proposed

Standard), RFC Editor, Fremont, CA, USA, Oct. 2011. [Online]. Available: https://www.rfc-editor.org/rfc/rfc6396.txt
[52] Internet Society, "MANRS for network operators," Internet Society, 2022. [Online]. Available: https://www.manrs.org/network-operators/
[53] J. Postel, "Internet Protocol," RFC 791 (Internet Standard), RFC Editor, Fremont, CA, USA, Sep. 1981, updated by RFCs 1349, 2474, 6864. [Online]. Available: https://www.rfc-editor.org/rfc/rfc791.txt
[54] J. Abley, B. Dickson, W. Kumari, and G. Michaelson, "AS112 Redirection Using DNAME," RFC 7535 (Informational), RFC Editor, Fremont, CA, USA, May 2015. [Online]. Available: https://www.rfc-editor.org/rfc/rfc7535.txt
[55] J. Abley and W. Sotomayor, "AS112 Nameserver Operations," RFC 7534 (Informational), RFC Editor, Fremont, CA, USA, May 2015. [Online]. Available: https://www.rfc-editor.org/rfc/rfc7534.txt
[56] M. Ziv, L. Izhikevich, K. Ruth, K. Izhikevich, and Z. Durumeric, "ASdb: A System for Classifying Owners of Autonomous Systems," in *ACM Internet Measurement Conference*, 2021, pp. 703–719.
[57] "AS to organizations mappings." [Online]. Available: https://catalog.caida.org/dataset/as_organizations
[58] CAIDA, "ASRank," CAIDA. [Online]. Available: https://asrank.caida.org/
[59] IIJ, "Internet health report," 2024. [Online]. Available: https://ihr.iijlab.net/ihr/hegemony/countries/JP
[60] "Country stats for last year of data," https://spoofer.caida.org/country_stats.php, accessed: 2024-05-13.
[61] Y. He, M. Faloutsos, S. Krishnamurthy, and B. Huffaker, "On routing asymmetry in the Internet," in *IEEE Global Telecommunications Conference*, vol. 2, 2005, pp. 6–.
[62] W. De Vries, J. J. Santanna, A. Sperotto, and A. Pras, "How asymmetric is the Internet? A Study to Support the use of Traceroute," in *IFIP International Conference on Autonomous Infrastructure, Management and Security*. Springer, 2015, pp. 113–125.
[63] K. Vermeulen, E. Gurmericliler, I. Cunha, D. Choffnes, and E. Katz-Bassett, "Internet scale reverse traceroute," in *ACM Internet Measurement Conference*, 2022, p. 694–715.
[64] P. Sermpezis, L. Prehn, S. Kostoglou, M. Flores, A. Vakali, and E. Aben, "Bias in Internet Measurement Platforms," in *2023 7th Network Traffic Measurement and Analysis Conference (TMA)*. IEEE, 2023, pp. 1–10.
[65] X. Deng and D. Loguinov, "BGP-Multipath Routing in the Internet," *IEEE/ACM Transactions on Networking*, vol. 29, no. 4, pp. 1708–1720, 2021.
[66] Q. B. Lone, M. Korczyński, C. H. Gañán, and M. J. G. van Eeten, "SAVing the internet: Explaining the adoption of source address validation by internet service providers," in *Workshop on the Economics of Information Security*, 2020, pp. 1–17.
[67] M. Luckie, R. Beverly, R. Koga, K. Keys, J. A. Kroll, and k. claffy, "Network hygiene, incentives, and regulation: Deployment of source address validation in the Internet," in *ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2019, pp. 2345–2361.
[68] A. Alaraj, K. Bock, D. Levin, and E. Wustrow, "A global measurement of routing loops on the internet," in *International Conference on Passive and Active Network Measurement*. Springer, 2023, pp. 373–399.
[69] S. Bellovin, R. Bush, T. Griffin, and J. Rexford, "Slowing routing table growth by filtering based on address allocation policies," 2001. [Online]. Available: https://www.cs.princeton.edu/~jrex/papers/filter.pdf
[70] N. Feamster, J. Jung, and H. Balakrishnan, "An Empirical Study of "Bogon" Route Advertisements," *ACM SIGCOMM Computer Communication Review*, vol. 35, no. 1, pp. 63–70, 2005.
[71] RIPE NCC, "RIPE Atlas: A Global Internet Measurement Network," *Internet Protocol Journal*, vol. 18, no. 3, pp. 2–26, 2015.
[72] M. A. Canbaz, J. Thom, and M. H. Gunes, "Comparative analysis of internet topology datasets," in *IEEE Conference on Computer Communications Workshops*, 2017, pp. 635–640.
[73] L. Salamatian, T. Arnold, Í. Cunha, J. Zhu, Y. Zhang, E. Katz-Bassett, and M. Calder, "Who Squats IPv4 Addresses?" *ACM SIGCOMM Computer Communication Review*, vol. 53, no. 1, pp. 48–72, 2023.
[74] T. Dai and H. Shulman, "Smap: Internet-wide scanning for spoofing," in *Annual Computer Security Applications Conference*, 2021, pp. 1039–1050.
[75] O. Fonseca, Í. Cunha, E. Fazzion, W. Meira, B. A. da Silva, R. A. Ferreira, and E. Katz-Bassett, "Identifying networks vulnerable to IP spoofing," *IEEE Transactions on Network and Service Management*, vol. 18, no. 3, pp. 3170–3183, 2021.