Internet Engineering Task Force
Internet-Draft
Intended status: Standards Track
Expires: 4 September 2025

S. Wang
Zhongguancun Laboratory
D. Li
Tsinghua University
L. Chen
R. Li
Zhongguancun Laboratory
H. Lin
Tsinghua University

3 March 2025

Source Address Validation Deployment Status
draft-wang-sav-deployment-status-00

Abstract

This document provides a summary of methods for measuring the deployment status of source address validation, with an overview of its deployment status. It reviews various methods for measuring outbound and/or inbound source address validation, including established tools like CAIDA Spoofer, as well as recently proposed remote measurement methods. By combining results from these different methods, the document offers a comprehensive overview of the status of source address validation deployment across the Internet.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 4 September 2025.

Copyright Notice

Table of Contents

1. Introduction

IP spoofing, sending packets with source addresses that do not belong to the sending host, is one of the long-standing security threats in the Internet. Source address validation (SAV) is important for protecting networks from IP spoofing attacks. Several techniques have been proposed to validate the source address of traffic, including Access Control List (ACL), unicast Reverse Path Forwarding (uRPF), and Virtual routing and forwarding (VRF) table. SAV can be categorized into two types: outbound SAV (OSAV) and inbound SAV (ISAV). OSAV discards spoofed packets originating from within a network and destined for external destinations, while ISAV focuses on filtering spoofed packets arriving from external sources to the

network.

The MANRS initiative considers IP spoofing as one of the most common routing threats, and defines a recommended action to mitigate spoofing traffic [manrs], encouraging network operators to implement SAV for their own infrastructure and end users, and for any Single-Homed Stub Customer Networks. However, as a recommended action, not all MANRS members follow this action to implement SAV for their networks, and only 1.6% of all routed ASes participate in MANRS. As a result, there is a lack of comprehensive knowledge regarding the current status of SAV deployment across the Internet community.

This document aims to provide a comprehensive view about SAV deployment in the Internet. The topics discussed in this document are organized into three main sections.

* Section 3 summarizes methods for measuring the deployment of OSAV.

* Section 4 summarizes methods for measuring the deployment of ISAV.

* Section 5 describes and analyzes the SAV deployment based on the measurement results derived from these methods.

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. Terminology

Spoofed Packet:
: A packet with forged source IP address. That is, the source IP address of the packet is not the legitimate IP address assigned to the sender.

Outbound Spoofing:
: The behavior where a network does not discard spoofed packets sent from the network to the outside.

Inbound Spoofing:
: The behavior where a network does not discard spoofed packets sent from the outside to the network.

Outbound Source Address Validation (OSAV):

The mechanism that discards spoofed packets sent from a network to the outside of it.

Inbound Source Address Validation (ISAV):
The mechanism that discards spoofed packets sent from the outside of a network to it.

Filtering Granularity:
The granularity of source address validation. If filtering granularity is /8, the network allows packets to be sent with any address that belong to the same /8 prefix as its own address. However, if filtering granularity is /8, the network allows to receive packets with any address as the source address that belongs to a different /8 prefix than its own address.

Filtering Depth:
The deployment depth of souce address validation. If filtering depth is 3, the source address validation is deployed 3 hops away from the sender for OSAV.

Authoritative DNS Nameserver (ADNS):
A DNS server that holds the definitive records for a domain and responds to DNS queries for that domain.

3. Outbound Source Address Validation Measurement Methods

To measure whether a network deploys OSAV, a common idea of different methods is to send spoofed packets from the network inside, and observe whether the spoofed packets reach the outside of the network. The SAV research community has proposed 3 methods for measuring OSAV deployment, i.e., client-based method, proxy-based method and DNAT-based method.

3.1. Client-based Method

As shown in Figure 1, by deploying a measurement client on a host in the audited network, the client can actively generate and send spoofed packets to the outside of the audited network. Hence, it is easy to learn whether spoofed packets have reached the outside of the network. Also, the client can set the time-to-live (TTL) of spoofed packets incrementally, and thus the forwarding path of the spoofed packets can be learned in a way like traceroute.

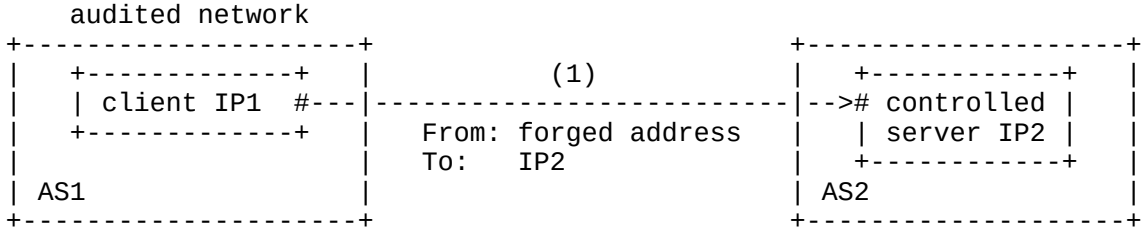

The client actively sends a set of spoofed packets to the controlled servers outside of the audited network.

Figure 1: An example of client-based OSAV measurement.

Benefiting from the controlbitly, a client can generate spoofed packets with arbitrary IP addresses as its source addresses. Hence, filtering granularity can be measured by observing which spoofed packets can reach outside of the audited network. Similarly, filtering depth can be measured by observing how far the spoofed packets reach.

The most famous client tool is the CAIDA Spoofer project [spoofer], which re-launched in 2015. A CAIDA Spoofer client sends various spoofed packets to a set of servers maintained by the project, and based on the spoofed packets received by the servers, the project is able to infer the filtering granularity of SAV on paths traversed by these packets. The CAIDA Spoofer project employs tracefilter to measure where a SAV mechanism is deployed. Speicifically, a client sends spoofed packets with specially crafted TTLs, and when the packet’s TTL expires, an ICMP TTL exceeded message will be sent to a controlled server. Based on the largest TTL among received ICMP messages, the project can infer the number of hops away from the client before spoofed packets are discarded.

The CAIDA Spoofer project relies on volunteers to spoof from many points in the network. If a volunteer installs the client within a Network Address Translation (NAT) network, CAIDA Spoofer will report the presence of a NAT device, and thus spoofed packets may be blocked by the NAT devices, rather than a SAV mechanism. Due to the wide deployment of NAT, though more than two thousands ASes were tested by the CAIDA Spoofer project in 2024, only about half of them were tested based on public IP addresses.

3.2. Proxy-based Method

[dns-proxy] relies on misbehaving DNS proxies to perform remote measurement of IP spoofing. As illustrated in Figure 2, the measurement conducter controls a scanner, a DNS authoritative nameserver, and a domain name, but does not have control over the audited network. The scanner first sends a DNS query for the domain name to a DNS proxy in the audited network, i.e., the destination IP address of the DNS query is the DNS proxy. However, due to the misbehaviors of the DNS proxy, it will forward the query to a DNS resolver without changing the source IP address of the query. In this way, the DNS proxy creates a spoofed packet whose source IP address does not belong to the audited network. If the spoofed packet is not discared along the path, the DNS resolver will communicate with the controlled authoritative nameserver to resolve the domain name. Finally, the DNS resolver will directly respond to the scanner, since the source IP address of the DNS query received by the DNS resolver is the scanner. Hence, if the scanner receives a DNS response whose source address is different from the destination address of the DNS query, the network is considered to have no OSAV deployment.

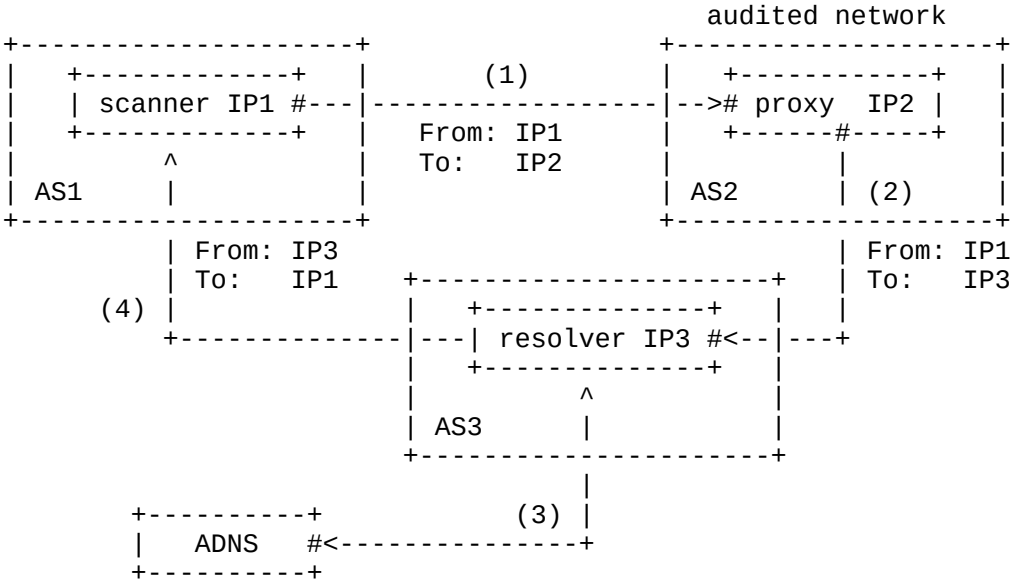

The scanner sends a DNS query with IP1 as the source to the DNS proxy (IP2). The proxy forwards the query to the DNS resolver, with the source IP address remaining as IP1. The resolver resolves the domain name using the authoritative name servers and responds directly to the scanner.

Figure 2: An example of proxy-based OSAV measurement.

Note that the IP address of the DNS proxy is also encoded into the domain name before sending to the DNS proxy, so that a DNS response can be matched with the corresponding DNS query. In addition, there is no need to find misbehaving DNS proxies before sending DNS queries to them. Instead, we can send DNS queries directly to all the routable address space one by one. If the destination address of a DNS query is used by a misbehaving DNS proxy, the scanner will receive a DNS response with an unexpected source address.

Proxy-based method can efficiently identify networks that do not deploy OSAV in a remote manner, but fails to identify networks that deploy OSAV. This is because, if OSAV is deployed in the audited network, the scanner will receive no DNS response, which is indistinguishable from the absence of a DNS proxy in the audited network.

3.3. DNAT-based Method

[DNAT] improves the proxy-based method by extending more than DNS protocol, identifying the deployment location of OSAV, and identifying the filtering granularity. Specifically, [DNAT] first figures out that the root cause of misbehaving DNS proxies is misconfigured destination NAT (DNAT) devices. As shown in Figure 3, when a packet matches DNAT rules, the DNAT device changes the packet’s destination to a preset address, while leaving the source address unchanged. Hence, to improve measurement coverage, DNAT-based method can also utilize other protocols, such as Network Time Protocol (NTP) and TCP protocol, to trigger the audited network into sending spoofed packets.

DNAT-based method identifies the filtering depth in a similar way to tracefilter. As DNAT devices do not reset the TTL field when forwarding packets, the fowarding path taken by spoofed packets can be learned by gradually incrementing the initial TTL values in original packets. The last responsive hop is considered as the position where filtering happens.

To identify the filtering granularity, the scanner sends multiple original packets with various source IP addresses. By using addresses adjacent to IP2 as the source addresses, the DNAT device will send spoofed packets with these addresses. Hence, packets that use forged addresses within the filtering granularity as source address will reach the receiver IP3.

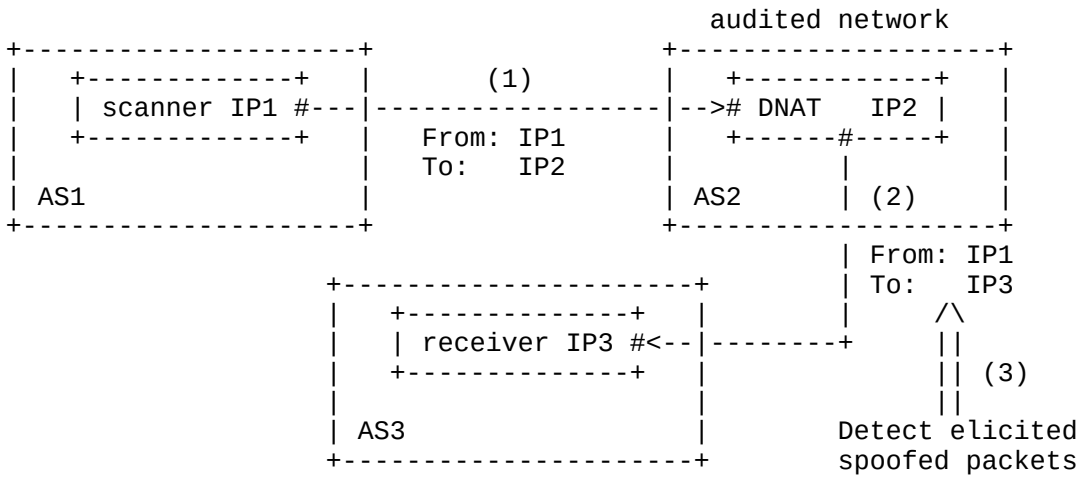

The scanner sends a packet sourced with IP1 to the DNAT device (IP2). The packet will elicit a spoofed packet sourced with IP1 and destined to IP3.

Figure 3: An example of DNAT-based OSAV measurement.

4. Inbound Source Address Validation Measurement Methods

The core idea of measuring whether a network deploys ISAV is to first send some spoofed packets to the target network and then observe whether the spoofed packets arrive inside of the target network. Since ISAV measurement does not require hosts in the audited network to generate spoofed packets, it is easier to measure ISAV deployment than OSAV. The SAV research community has proposed 5 methods for measuring OSAV deployment, i.e., client-based method, resolver-based method, ICMPv6-based method, IPID-based method and PMTUD-based method.

4.1. Client-based Method

As shown in Figure 4, by deploying a measurement client on a host in the audited network, client-based method can use a controlled server to send a spoofed packet to the client. The spoofed packets use an IP addresses adjacent to IP2 as its source IP addresses. If the client receives the spoofed packet, then the audited network has not deployed ISAV. Otherwise, the audited network has deployed ISAV.

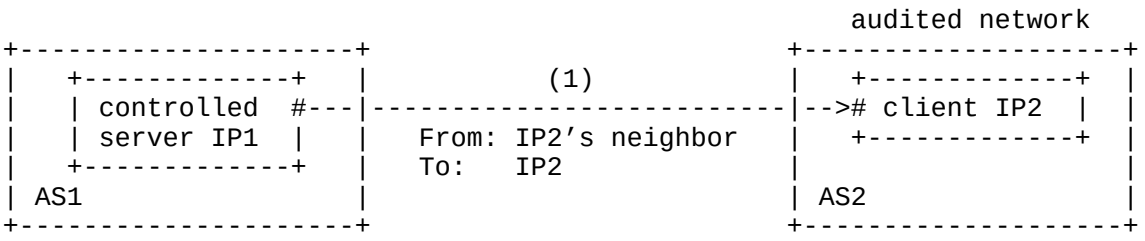

The controlled server sends a spoofed packet to the client, and then client reports whether it has received the spoofed packets.

Figure 4: An example of client-based ISAV measurement.

The CAIDA Spoofer project [spoofer] also supports ISAV measurements, which, like OSAV measurements, rely on volunteers. When volunteers install the CAIDA Spoofer client, both ISAV and OSAV measurements are performed on the audit network. However, if the client is installed within a NAT network, it becomes inaccessible from outside the network, even without spoofed addresses. As a result, client-based methods cannot measure ISAV deployments in this case.

4.2. Resolver-based Method

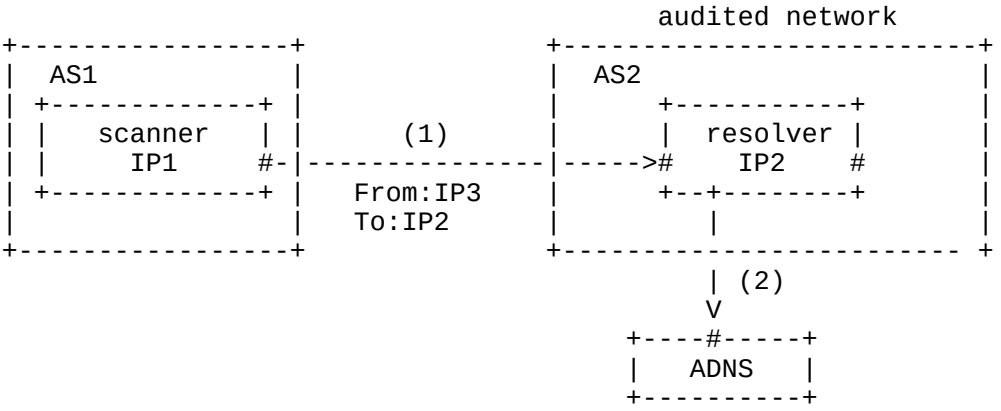

Figure 5: An example of resolver-based ISAV measurement.

Figure 5 shows an example of resolver-based ISAV measurement [dns-resolver]. The scanner in AS1 sends a DNS query with a forged IP address IP3, which belongs to the audited network (AS2), to a DNS resolver in AS2. If the audited network does not deploy ISAV, the DNS resolver will receive the spoofed DNS query. Next, the DNS resolver will send another DNS query to our controlled ADNS for resolution. Therefore, if the controlled ADNS receives the DNS query from the DNS resolver in the audited network, the audited network AS2 does not filter the spoofed packets.

However, if the controlled ADNS does not receive the DNS query, we can not assume that the audited network filters the spoofed packets, since there may be no DNS resolver in the audited network. To futher identify networks that filter inbound spoofing traffic, we send a non-spoofed DNS query from the scanner to the same target IP address. If the controlled ADNS receives a DNS query triggered by the non-spoofed DNS query, a DNS resolver exists in the audited network. As a result, if the DNS resolver does not resolve the spoofed query, we can conclude that the audited network deploy ISAV.

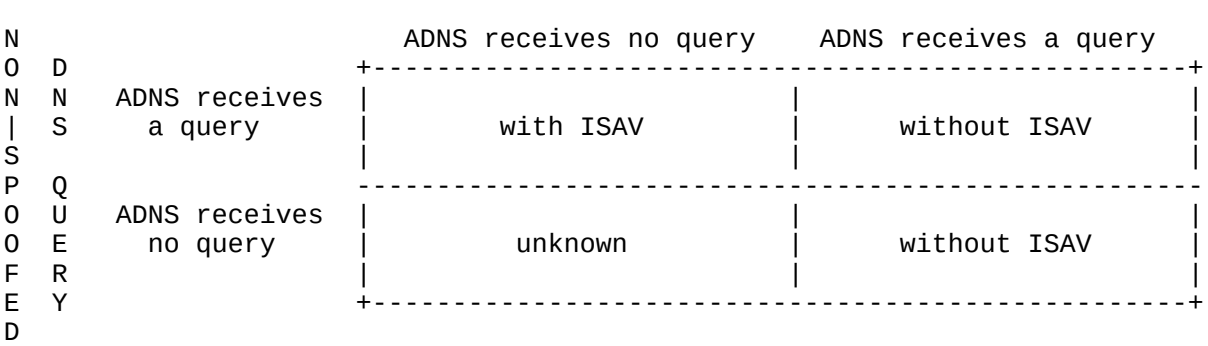

Figure 6: Classification of results based on DNS resolvers.

As illustrated in Figure 6, there are four cases when combining spoofed DNS query and non-spoofed DNS query.

* First, the ADNS receives DNS queries in both spoofing scan and non-spoofing scan, suggesting that the audited network does not deploy ISAV, and the DNS resolver is open.

* Second, the ADNS receives the DNS query only in spoofing scan, suggesting that the audited network does not deploy ISAV, and the DNS resolver is closed.

* Third, the ADNS receives the DNS query only in non-spoofing scan, suggesting that the audited network deploys ISAV.

* Fourth, the ADNS does not receive any DNS query. This suggests that no DNS resolver in the audited network can be utilized to measure ISAV deployment.

4.3. ICMPv6-based Method

As suggested by [RFC4443], in order to limit the bandwidth and forwarding costs incurred by originating ICMPv6 error messages, an IPv6 node MUST limit the rate of ICMPv6 error messages it originates. This provides an opportunity to infer whether the spoofed packets arrive inside of the audited network based on the state of ICMPv6 rate limiting. Since most of IPv6 addresses are inactive, an ICMP error message will be fed back from Customer Premises Equipment (CPE) devices when we send an ICMP echo request to a random IP address in the audited network. If the CPE device limits the rate of ICMPv6 error messages it originates, it can be utilized as a vantage point (VP).

Figure 7 illustrates the ICMPv6-based measurement method [ICMPv6]. We have a local scanner P1 in AS1, and AS2 is the audited network. Three rounds of testing with N and N+M ICMP echo requests packets are conducted in the measurement, and thus three values rcv1, rcv2, and rcv3 can be obtained respectively. Based on this, we can infer whether the audited network filters the spoofed packets by comparing rcv1, rcv2, and rcv3.

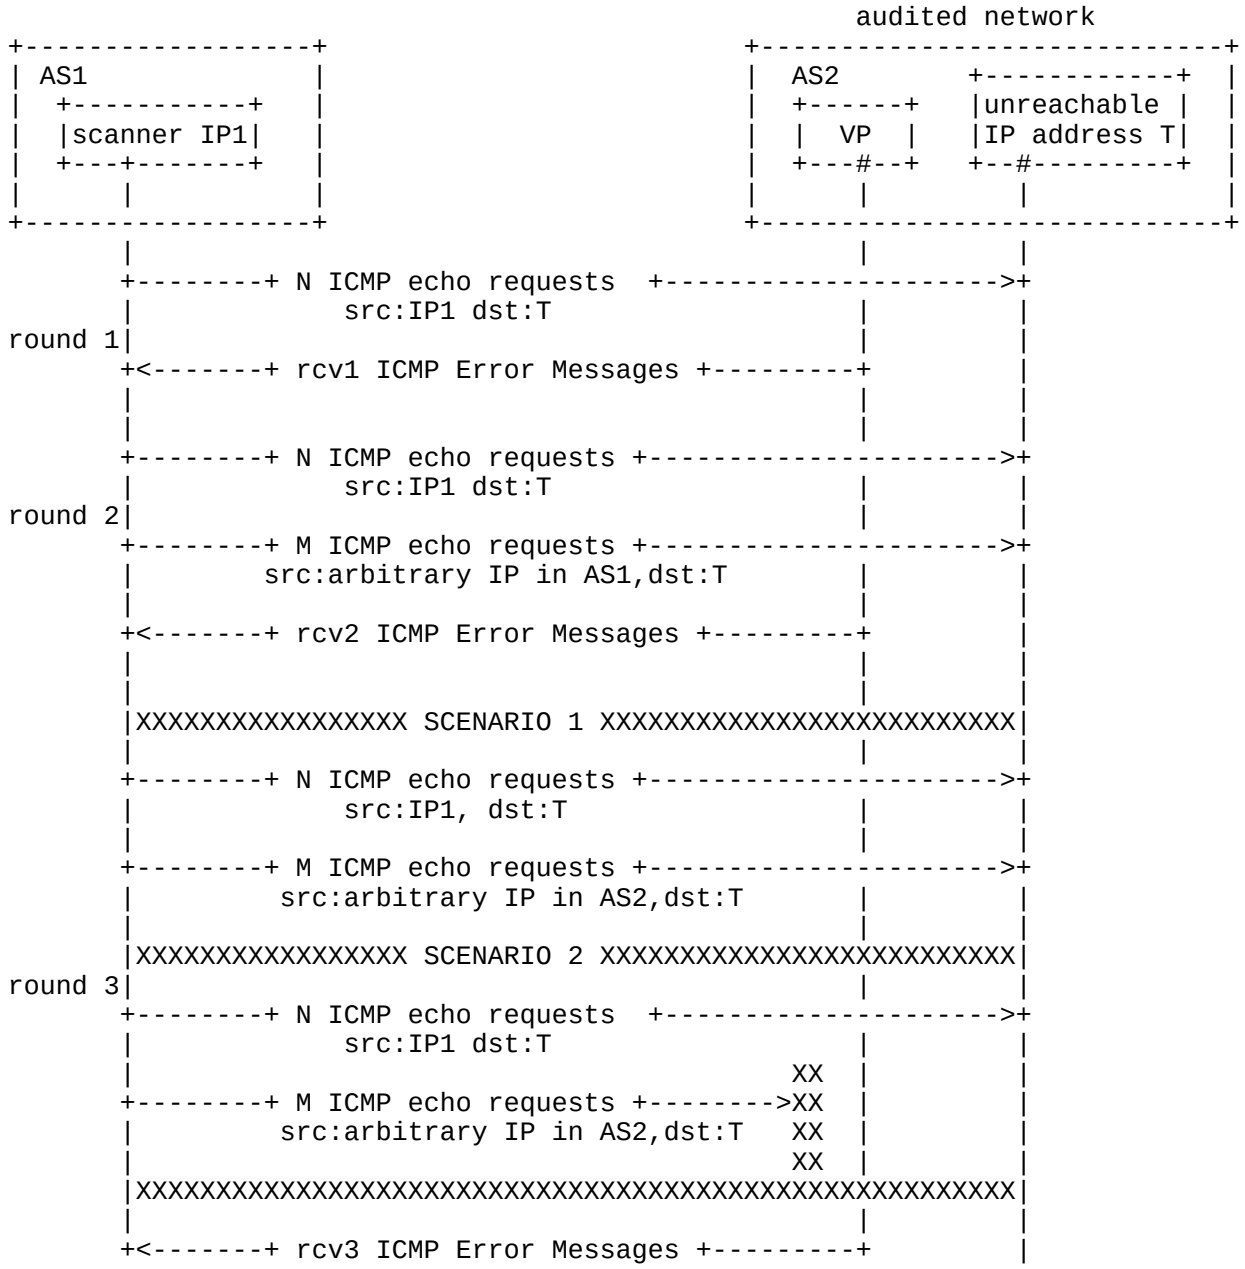

Figure 7: An example of ICMPv6-based ISAV measurement.

As illustrated in Figure 7, in the first round test, N ICMP echo requests are sent to a target with inactive IPv6 address in the audited network, and then VP will resposnd with rcv1 ICMP error messages to the scanner. In the second round test, besides the same N probe packets, extra M ICMP echo requests with forged source address that belongs to AS1 (noise packets) are sent to the target simultaneously. The number of ICMP error messages in the second round test are defined as rcv2. Similar to the second round test, in the third round test, M ICMP echo requests with forged source address that belongs to AS2 (spoofed packets) are sent to the target. The number of ICMP error messages in the third round test are defined as rcv3.

Comparing rcv1 and rcv3, if rcv1 > rcv3, it can be considered that the spoofed packets are not filtered in the third round test, suggesting that the audited network allows inbound spoofing. Comparing rcv2 and rcv3, if rcv2 < rcv3, it can be inferred that the target network has filtered the spoofed packet in the third round test, and thus is able to filter inbound spoofing traffic. The ability of filtering inbound spoofing traffic can be inferred according to the following rules.

* If rcv3 < a*rcv1, then the network allow inbound spoofing;

* Else if rcv2 < a*rcv3, then the network does not allow inbound spoofing;

* Else, the ability of filtering inbound spoofing traffic cannot be determined.

where a is a factor to avoid potential interference from fast-changing network environments, satisfying 0 < a < 1.

4.4. IPID-based Method

The core observation of using IPID to measure ISAV is that the globally incremental IPID value leaks information about traffic reaching the server[SMap]. Given a server in the audited network with a globally incremental IPID, the scanner samples the IPID value using its own IP address by sending packets to the server and receiving responses. Then, the scanner sends a set of packets to the server using a spoofed IP address that belongs to the audited network, i.e., an IP address adjacent to IP2. Afterward, the scanner sends another packet using its IP address to probe the IPID value again. If the spoofed packets reached the server, they would have incremented the server’s IPID counter. As a result, this increment can be inferred during the second IPID probe from the scanner’s IP address.

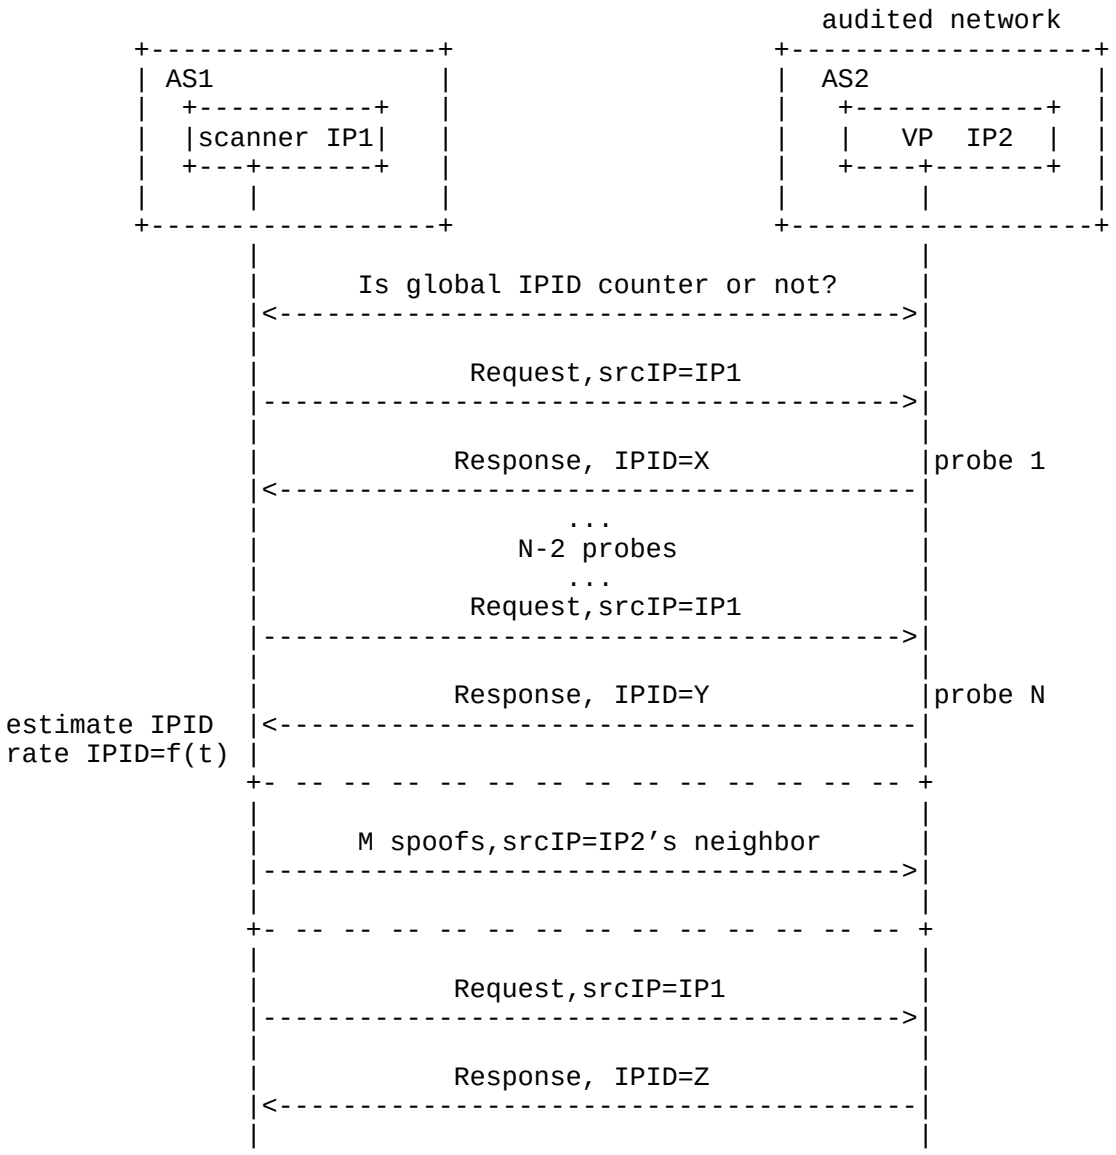

Figure 8: An example of IPID-based ISAV measurement.

Figure 8 illustrates the measurement process of ISAV based on global IPID. First, the scanner measures the current IPID value and the rate of IPID increments. Ordinary Least Squares (OLS) linear regression can be used to estimate the relationship between the IPID and the timestamp t: IPID = a*t + b + ε , ε N (0, σ^2). Next, N probes are sent to the VP. With these N probes, the parameters a, b,

and σ can be estimated using the OLS method. Then, a group of M = 6 * σ packets with a spoofed source IP address are sent to the audited network. Finally, the IPID value Z from the VP is sampled by using IP1 as source address, while the IPID value W at that moment can be estimated using the linear regression model. If the M spoofed packets are filtered, according to the 3-sigma rule, there is a 99.73% probability that the following condition holds: W - 3 * σ Z W + 3 * σ . If the spoofed packets are not filtered, meaning the audited network has not deployed ISAV, the IPID counter will increase by M. In this case, Z > W + 3 * σ , or equivalently, Z > W + M/2.

4.5. PMTUD-based Method

The core idea of the Path MTU Discovery (PMTUD) method is to send ICMP Packet Too Big (PTB) messages with a spoofed source IP address that belongs to the audited network [SMap]. The real IP address of the scanner is embedded in the first 8 bytes of the ICMP payload. If the network does not deploy ISAV, the server will receive the PMTUD message and reduce the MTU for the IP address specified in the first 8 bytes of the ICMP payload. First, probe the MTU of the service in the audited network. Then, send an ICMP PTB message from a spoofed IP address. If the packet reaches the service, it will reduce the MTU for the scanner's IP address. This reduction will be identified in the next packet received from the service, indicating that the audited network does not deploy ISAV.

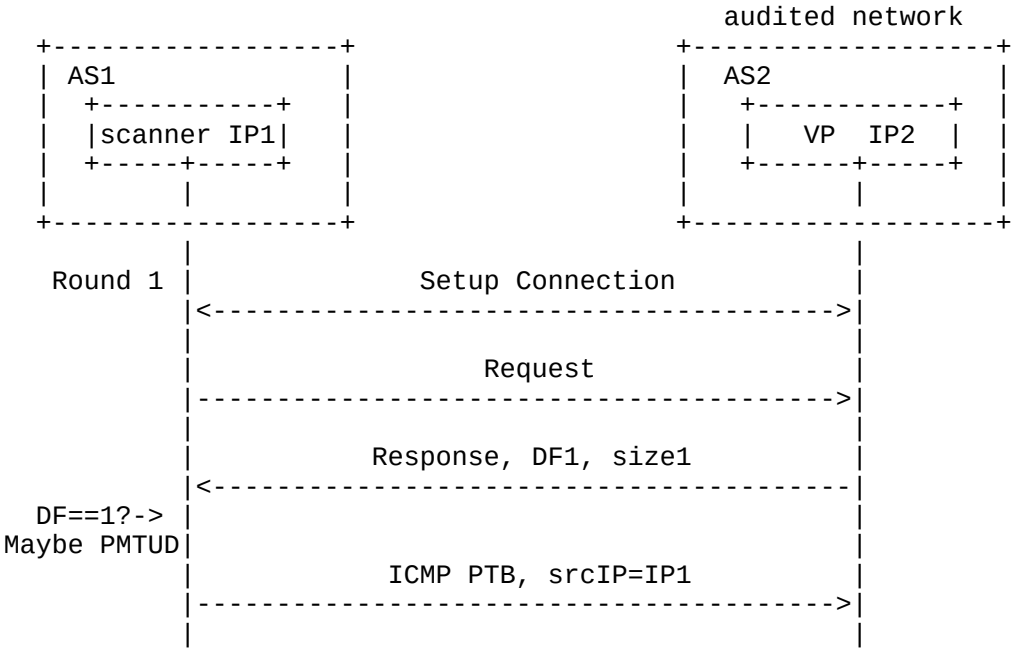

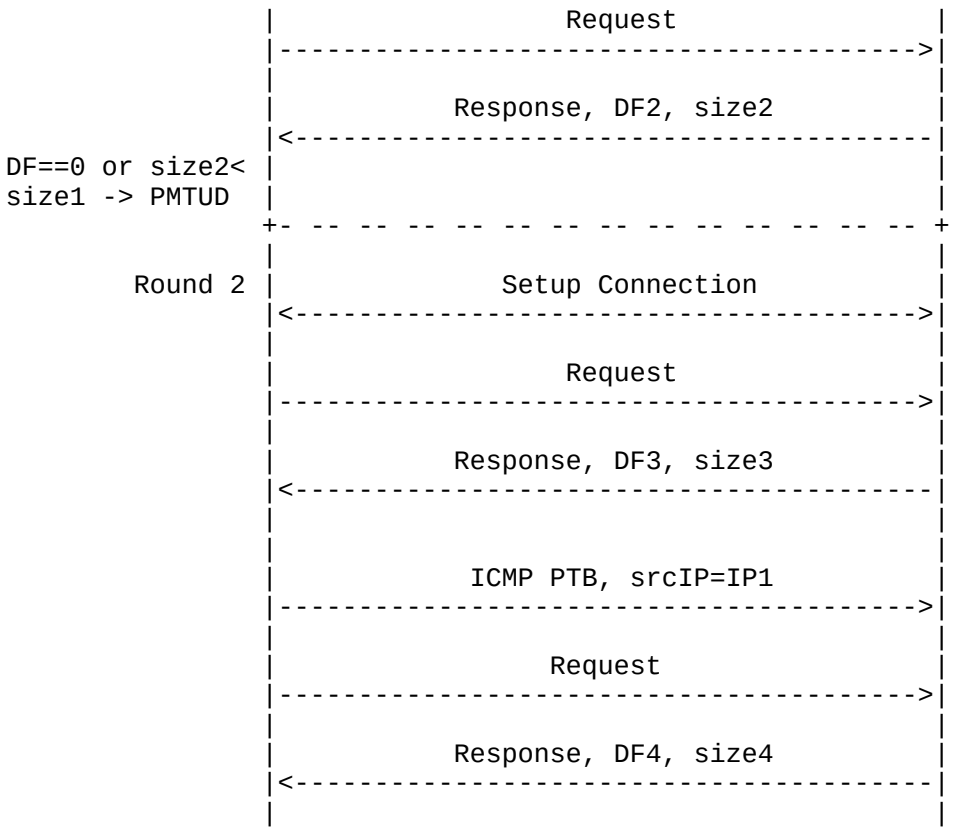

Figure 9: An example of PMTUD-based ISAV measurement.

Figure 9 illustrates the measurement process of ISAV based on PMTUD. First, establish a TCP connection with the server in the audited network. Then, send Request1 and receive Response1. If the DF (Don't Fragment) bit is not set, the server does not support PMTUD. Otherwise, send an ICMP PTB message with a smaller MTU. Next, issue another request and receive Response2. If DF1 == 1 and (DF2 == 0 or size2 size1), the server supports PMTUD. Now, proceed to test whether ISAV is deployed. Use the neighbor's IP address of the server as the source IP address to spoof an ICMP PTB with the smallest MTU. After that, issue another request. If the following condition is observed, the server is not protected by ISAV: size4 size3 or (DF3 == 1 and DF4 == 0).

5. Deployment Status

5.1. Global Picture

In February 2025, we used the above methods to measure SAV deployment in the Internet. As shown in Figure 11 and Figure 10, 67.4% of IPv4 and 72.8% of IPv6 ASes lacked any ISAV deployment. Partial deployment was observed in 30.2% of IPv4 and 23.1% of IPv6 ASes, suggesting that these ASes deploy ISAV at their access networks.

Category	IPv4 ASes	IPv6 ASes
Deployed	1,157 (2.5%)	372 (4.0%)
Not Deployed	31,817 (67.4%)	6,747 (72.8%)
Partially Deployed	14,235 (30.2%)	2,143 (23.1%)

Figure 10: ISAV deployment status across IPv4 ASes and IPv6 ASes.

Category	IPv4 /24 Prefixes	IPv6 /48 Prefixes
Deployed	222,362 (13.1%)	47,704 (8.1%)
Not Deployed	1,390,206 (82.0%)	404,629 (68.5%)
Partially Deployed	83,460 (4.9%)	138,693 (23.5%)

Figure 11: ISAV deployment status across IPv4 /24 and IPv6 /48 prefixes.

Figure 13 and Figure 12 show OSAV deployment disparities between IPv4 and IPv6 networks. Only 14.8% of IPv4 ASes and 17.8% of IPv4 /24 prefixes demonstrate complete OSAV deployment. In contrast, IPv6 compliance is significantly higher than IPv4.

Category	IPv4 ASes	IPv6 ASes
Deployed	409 (14.8%)	318 (71.6%)
Undeployed	2,200 (79.6%)	81 (18.2%)
Partially Deployed	155 (5.6%)	45 (10.1%)

Figure 12: OSAV deployment status across IPv4 and IPv6 ASes.

Category	IPv4 /24 Prefixes	IPv6 /48 Prefixes
Deployed	1,402 (17.8%)	679 (80.9%)
Undeployed	6,335 (80.4%)	130 (15.5%)
Partially Deployed	140 (1.8%)	30 (3.6%)

Figure 13: OSAV deployment status across IPv4 /24 and IPv6 /48 prefixes.

Figure 14 shows OSAV deployment granularity patterns. The prefix length of /20-/24 dominates deployment (55.52%), as these prefix lengths correspond to standard IPv4 allocation units for ASes. This pattern suggests OSAV is predominantly deployed at AS border interfaces.

Range	Percentage
/ 8	0.13 %
/ 9	0.26 %
/10	0.53 %
/11	0.13 %
/12	0.26 %
/13	0.66 %
/14	0.79 %
/15	0.53 %
/16	3.95 %
/17	4.74 %
/18	3.29 %
/19	5.53 %
/20	6.97 %
/21	8.55 %
/22	23.95 %
/23	7.76 %
/24	8.29 %
/25	2.24 %
/26	2.63 %
/27	3.95 %
/28	3.29 %
/29	5.79 %
/30	3.42 %
/31	2.37 %

Figure 14: OSAV filtering granularity in IPv4 networks.

Figure 15 shows the filtering granularity of ISAV, with 41.66% of networks implementing spoofing filters at /29-/30 granularity (per IETF BCP38 recommendations). This suggests ISAV is predominantly deployed in access networks.

Range	Percentage
/ 8	0.17 %
/ 9	1.99 %
/10	6.07 %
/11	4.48 %
/12	4.94 %
/13	3.50 %
/14	3.99 %
/15	5.78 %
/16	2.17 %
/17	3.27 %
/18	2.76 %
/19	2.43 %
/20	1.84 %
/21	3.25 %
/22	1.73 %
/23	3.24 %
/24	1.55 %
/25	0.97 %
/26	1.02 %
/27	1.35 %
/28	1.85 %
/29	22.94 %
/30	18.72 %

Figure 15: ISAV filtering granularity in IPv4 networks.

Figure 16 characterizes OSAV filtering depth measured by the DNAT-based method, where 91.52% of deployment are within 2 IP hops from the endpoints - with complete absence beyond 10 hops.

Hop	Percentage
1	66.01 %
2	25.51 %
3	4.58 %
4	2.46 %
5	1.03 %
6	0.14 %
7	0.00 %
8	0.21 %
9	0.07 %
10	0.00 %

Figure 16: OSAV filtering depth in IPv4 networks.

5.2. Deployment in Countries

The SAV deployment in the global Internet is shown in Figure 18 and Figure 17. Analysis of regions with sufficient data reveals distinct deployment patterns: China, South Korea, Germany, and France demonstrate higher OSAV deployment ratios, while Russia, Brazil, and India show lower OSAV deployment ratios. Notably, ISAV deployment remains limited in most regions, with South Korea, Poland, and Egypt emerging as exceptional cases exhibiting more advanced ISAV deployment.

Country	OSAV Tested Prefixes	OSAV Deployment Ratio
CN	376	76.3%
KR	58	75.9%
FR	12	75.0%
DE	16	68.8%
US	300	42.7%
NL	18	33.3%
PL	70	32.9%
CA	117	32.5%
GB	28	32.1%
AU	11	27.3%
IT	116	23.3%
TW	19	21.1%
EG	56	19.6%
ID	490	17.8%
JP	17	17.6%
MX	36	13.9%
ES	38	10.5%
RU	75	9.3%
BR	2,575	7.3%
IN	1,430	5.5%

Figure 17: OSAV deployment among countries/regions.

Country	ISAV Tested Prefixes	ISAV Deployment Ratio
KR	71,934	44.8%
TW	22,523	42.0%
PL	17,880	40.5%
EG	16,806	37.3%
FR	35,220	19.4%
DE	49,956	14.4%
ES	15,018	16.2%
BR	47,874	11.8%
US	562,655	10.2%
RU	56,084	10.2%
AU	21,023	8.3%
NL	19,803	8.3%
CA	23,801	7.2%
GB	31,271	6.9%
JP	67,173	6.2%
IT	30,357	5.7%
CN	211,539	4.8%
ID	18,845	4.6%
IN	30,569	4.1%
MX	17,665	3.4%

Figure 18: ISAV deployment among countries/regions.

5.3. Comparison between ISAV and OSAV

Figure 19 and Figure 20 show the deployment status of ISAV and OSAV across ASes. Our measurements focus on ISP ASes, revealing significant disparities:

ISAV: China Telecom (AS4134), China Unicom (AS4837), AT&T (AS7018), and Verizon (AS701) exhibit low deployment rates. In contrast, Korea Telecom (AS4766), Comcast (AS7922), Charter (AS20115), and Chungwa Telecom (AS3462) demonstrate significantly higher ISAV deployment.

OSAV: China Telecom (AS4134), China Unicom (AS4837), and Korea Telecom (AS4766) achieve over 90% OSAV deployment across their /24 networks.

ASN	ISAV Tested Prefixes	ISAV Deployment Ratio
3462	12,752	70.1%
4766	37,667	60.8%
20115	13,505	40.1%
7922	29,403	22.9%
8075	22,415	10.0%
209	11,435	7.9%
12389	12,288	5.0%
3320	14,684	4.7%
4134	69,625	4.4%
4837	48,749	4.0%
7018	31,888	3.3%
4713	14,727	3.2%
16509	48,563	3.2%
45090	11,168	3.0%
3269	14,181	3.1%
701	15,694	2.2%
17676	11,702	1.8%
8151	11,996	1.6%
749	70,399	1.2%
36947	11,339	0.4%

Figure 19: ISAV deployment ratio of ASes.

ASN	OSAV Tested Prefixes	OSAV Deployment Ratio
272122	160	100.0%
14061	37	100.0%
4766	36	97.2%
4134	232	92.7%
4837	48	81.2%
17995	71	74.6%
15924	56	26.8%
8452	49	12.2%
38758	47	4.3%
150008	102	0.0%
34984	78	0.0%
52468	66	0.0%
395582	64	0.0%
58659	55	0.0%
23688	43	0.0%
18229	40	0.0%
52444	36	0.0%
133676	34	0.0%
23923	33	0.0%
18002	33	0.0%

Figure 20: OSAV deployment ratio of ASes.

We find a positive correlation between the deployment of OSAV and ISAV. That is, 10.9% of ASes that deploy ISAV also deploy OSAV, while only 5.9% of ASes without ISAV deploy OSAV. Similarly, 36.0% of ASes that deploy OSAV also deploy ISAV, while only 22.6% of ASes without OSAV deploy ISAV.

5.4. Impact of MANRS

To understand the impact of MANRS on SAV deployment, we compare SAV deployment ratios between MANRS and non-MANRS networks, including both full and partial deployments.

The analysis reveals MANRS networks demonstrate superior SAV deployment: 29.1% in MANRS networks versus 19.6% in non-MANRS networks for OSAV, and 73.3% vs. 56.7% for ISAV. These results indicate that although anti-spoofing is a recommended action, MANRS participation improves SAV deployment across network configurations.

	OSAV Deployment Ratio	ISAV Deployment Ratio
MANRS	29.1%	73.3%
Non-MANRS	19.6%	56.7%

Figure 21: The impact of MANRS on SAV deployment.

6. IANA Considerations

This document has no IANA requirements.

7. References

7.1. Normative References

[RFC4443] Conta, A., Deering, S., and M. Gupta, Ed., "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", STD 89, RFC 4443, DOI 10.17487/RFC4443, March 2006, <https://www.rfc-editor.org/rfc/rfc4443>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/rfc/rfc2119>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/rfc/rfc8174>.

7.2. Informative References

[spoofer] CAIDA, "Spoofer project", 2024, <https://spoofer.caida.org/>.

[manrs] MANRS, "MANRS Implementation Guide", 2024, <https://www.manrs.org/netops/guide/antispoofing/>.

[DNAT] "Remote Measurement of Outbound Source Address Validation Deployment", 2024, <https://datatracker.ietf.org/doc/draft-wang-savnet-remote-measurement-osav/>.

[dns-proxy]
Marc Kuhrer, Thomas Hupperich, Christian Rossow, and Thorsten Holz, Ruhr-University Bochum, "Exit from hell? Reducing the impact of amplification DDoS attacks", 2014, <https://www.usenix.org/system/files/conference/usenixsecurity14/sec14-paper-kuhrer.pdf>.

[dns-resolver]
Yevheniya Nosyk, Maciej Korczynski, Qasim Lone, Marcin Skwarek, Baptiste Jonglez, Andrzej Duda, "The Closed Resolver Project: Measuring the Deployment of Inbound Source Address Validation", 2023, <https://ieeexplore.ieee.org/document/10082958>.

[ICMPv6]
Long Pan, Jiahai Yang, Lin He, Zhiliang Wang, Leyao Nie, Guanglei Song, Yaozhong Liu, "Your Router is My Prober: Measuring IPv6 Networks via ICMP Rate Limiting Side Channels", 2023, <https://www.ndss-symposium.org/wp-content/uploads/2023/02/ndss2023_s49_paper.pdf>.

[SMap]
Tianxiang Dai, Haya Shulman, "Smap: Internet-wide Scanning for Spoofing", 2021, <https://dl.acm.org/doi/10.1145/3485832.3485917>.

Authors' Addresses

Shuai Wang
Zhongguancun Laboratory
Beijing
China
Email: wangshuai@zgclab.edu.cn

Dan Li
Tsinghua University
Beijing
China
Email: tolidan@tsinghua.edu.cn

Li Chen
Zhongguancun Laboratory
Beijing
China
Email: lichen@zgclab.edu.cn

Ruifeng Li
Zhongguancun Laboratory
Beijing
China
Email: lirf@zgclab.edu.cn

He Lin
Tsinghua University
Beijing
China
Email: he-lin@tsinghua.edu.cn