Grey Rhino Warning: IPv6 is Becoming Fertile Ground for Reflection Amplification Attacks

Ling Hu, Tao Yang, Yu Pang, Bingnan Hou, Zhiping Cai, Bo Yu

National University of Defense Technology, College of Computer Science and Technology

***Abstract*—Distributed Denial-of-Service (DDoS) attacks represent a cost-effective and potent threat to network stability. While extensively studied in IPv4 networks, DDoS implications in IPv6 remain underexplored. The vast IPv6 address space renders brute-force scanning and amplifier testing for all active addresses impractical. Innovatively, this work investigates AS-level vulnerabilities to reflection amplification attacks in IPv6.**

One prerequisite for amplification presence is that it is located in a vulnerable autonomous system (AS) without inbound source address validation (ISAV) deployment. Hence, the analysis focuses on two critical aspects: global detection of ISAV deployment and identification of amplifiers within vulnerable ASes. Specifically, we develop a methodology combining *ICMP Time Exceeded* mechanisms for ISAV detection, employ IPv6 address scanning for amplifier identification, and utilize dual vantage points for amplification verification.

Experimental results reveal that 4,460 ASes (61.36% of measured networks) lack ISAV deployment. Through scanning approximately 47M active addresses, we have identified reflection amplifiers in 3,507 ASes. The analysis demonstrates that current IPv6 networks are fertile grounds for reflection amplification attacks, alarming network security.

***Index Terms*—reflection amplification, ISAV, IPv6, security**

I. Introduction

Reflection amplification attacks remain a potent DDoS technique, where spoofed requests to vulnerable servers (e.g., DNS/NTP) generate up to $557\times$ amplified traffic [1]. The 2018 Memcached attack on GitHub [2] (peaking at 1.35 Tbps) demonstrated this threat. As IPv4 defenses improve, attackers are shifting focus to IPv6 networks [3]. With 35.67% global adoption [4], IPv6's vast address space provides abundant amplifiers, as seen in recent IPv6 DDoS attacks against Neustar [3]. This transition creates new security challenges in our new IPv6 landscape.

Key Observations. While the ongoing battle between attackers and defenders regarding reflection attacks in IPv4 has persisted for years [5]–[7], there is still a notable gap in comprehensive measurements of reflection attacks across the entire IPv6 network. Based on existing research, we can identify **three critical elements relevant to reflection amplification attacks** that apply to both IPv4 and IPv6.

1. **ASes, where attackers reside, lack OSAV**: The lack of outbound source address validation (OSAV) [8] in autonomous systems (ASes) where attackers are located allows these ASes to ignore the legitimacy of source addresses for outbound traffic, in turn facilitating source address spoofing.

2. **ASes housing amplifiers lack ISAV**: Amplifiers often exist within ASes that do not implement inbound source address validation (ISAV) [9]. This failure means these ASes do not verify whether source addresses of incoming traffic are legitimate according to network topology, allowing spoofed traffic to traverse their networks.

3. **Servers acting as amplifiers have vulnerable protocols**: Hosts utilized as amplifiers often provide weak services that are not properly configured, such as the Network Time Protocol (NTP) [10]. These corresponding protocols [1], [10] can generate response traffic that is orders of magnitude greater—often tens or even hundreds of times—than the initial query. Broad experiments [6], [7] have confirmed that hosts with these vulnerable protocols can be exploited as amplifiers to launch effective DDoS attacks.

Challenges. Given that the deployment detection of OSAV has been well-established [11], current vulnerability analysis in the IPv6 environment faces significant limitations primarily due to the challenges in detecting ISAV and identifying weak reflection points. This situation presents two main challenges.

Challenge 1: Efficient Detection of ISAV Deployment at a Global Scale. Traditional methods [12]–[14] rely on internal network volunteers or DNS resolvers, which are often impractical. IVANTAGE [15] utilizes a local probe through ICMP rate limiting mechanisms [16], but *requires sending a significant volume of packets in a short time*, which is inefficient, and can produce false positives by ignoring intermediate ASes. Finding a more efficient and accurate detection method under resource constraints is a key challenge.

Challenge 2: Identification of Amplifiers in Nearly Infinite Search Space. IPv6 has a 2^{96} times larger address space than IPv4, making exhaustive scanning impractical [17]. Furthermore, the conditions for forming amplifiers are even more stringent, as these hosts must be capable of reflecting TCP or UDP traffic, rather than *merely responding to ICMP traffic* [15]. The question of how to efficiently identify reflection points within this vast space remains an open challenge.

Our Paper. Testing every IPv6 address for amplification is infeasible, so we focus on identifying **AS-level vulnerabilities to reflection amplification attacks in IPv6**, targeting ASes without ISAV as potential amplifier hosts. We propose an efficient measurement scheme using ICMP probing with source address spoofing to detect ISAV, followed by IPv6 scanning to find live hosts and protocol testing to identify amplifiers. All measurements are conducted from two controlled vantage points, and detailed statistics and results are provided. Our code is available at https://gitee.com/ahaBCD/amplification.

Contributions. This paper makes significant contributions to understanding reflection amplification attacks within IPv6 networks by not only revealing potential scenarios for such attacks but also presenting a comprehensive analysis of their extensive characteristics.

- We analyze ISAV deployment at the AS level across the global IPv6 network. Our findings reveal that 4,460 ASes, 61.36% of all measured ASes, have yet to deploy ISAV. This comprehensive survey underscores the widespread existence of vulnerabilities within IPv6 networks.
- We assess the feasibility and success rate of executing reflection amplification attacks within identified ASes, using two controlled vantage points. It is shown that devices in 3,507 ASes are capable of reflecting and amplifying traffic, with amplification factors reaching up to $4,267\times$. This highlights the significant risks posed by inadequate security measures in these networks.
- Finally, we provide a detailed feature analysis of amplifiers across the IPv6 landscape and offer corresponding defense recommendations to enhance network security.

II. Background

In this section, we review the relevant concepts of reflection amplification and source address validation, concluding with an introduction to our measurement model.

A. Reflection Amplification

During reflection amplification, an attacker sends a request with a spoofed source address to a server acting as a reflector, enticing the server to respond to a victim. The scale and impact of the attack are determined by factors such as the number of reflectors, the volume of traffic, the duration of the attack, and so on. The amplification factor (AF) can be calculated as described in [1], [5]:

$$\begin{aligned} AF &= \#\ of\ amplifiers \cdot PAF \cdot SAF \cdot TCF\ , \\ PAF &= \frac{\#\ of\ response\ packets}{\#\ of\ query\ packets}\ , \\ SAF &= \frac{Response\ packet\ size}{Query\ packet\ size}\ , \\ TCF &= \frac{Time\ used\ for\ query}{Time\ used\ for\ response}\ . \end{aligned} \tag{1}$$

Traditional reflection amplification attacks primarily exploit protocols where the response size exceeds that of the query [1], [10]. Emerging attackers continually innovate, generating larger attack volumes with equal or fewer resources. For example, an attacker can make multiple spoofed *monlist* queries to an NTP server to expand the number of interactive clients in advance, prompting more response packets for one query [1]. Additionally, [18] successfully demonstrated the use of PSH packets to trigger reflection amplification attacks, thereby reducing the size of query packets. The method proposed in [5] allows attackers to send multiple DNS requests continuously to increase the time for query accumulating and minimize response time through *query aggregation* and *rapid response return* mechanisms. Moreover, attackers are progressively expanding their reflection types from traditional vulnerable servers to network middleboxes. Research by [18] indicates non-compliant traffic can induce middleboxes to return substantial traffic volumes without TCP three-way handshake required.

However, these studies predominantly focus on IPv4. This paper focuses on revealing the substantial potential for amplifiers within IPv6 and explores methods to increase amplification factor similarly.

B. Source Address Validation

Source address validation (SAV) is a critical technology for preventing source address spoofing by verifying the legality of packet source IP address, including inbound source address validation (ISAV) and outbound source address validation (OSAV) [8], [9]. OSAV prevents forwarding packets with spoofed source addresses from within its network, while ISAV restricts the entry of such packets into its network. OSAV deployment detection has matured, while ISAV remains a challenge [8], [9], [12], [13]. Existing ISAV deployment detection methods exhibit significant limitations: [19] is based on misconfigurations, [20] relies on passive traffic analysis, while Spoofer [12]–[14] demands probes within the network. Additionally, [9] requires comprehensive network scanning, which is impossible in IPv6, and [8] necessitates the deployment of open DNS resolvers.

The state-of-the-art IVANTAGE [15] relies on ICMP rate limiting, requiring a significant number of packets and risking misclassifying ASes without ISAV deployment as having it due to the ignorance of strict ISAV of the on-path routers. This paper introduces a more efficient and accurate method for assessing ISAV deployment within IPv6 networks, addressing these limitations.

C. Measurement Model

Our study employs minimally restrictive vantage points capable only of sending spoofed packets (without monitoring or modifying traffic) to ensure broad applicability in assessing AS-level vulnerabilities to reflection attacks. This approach is validated by CAIDA [11] showing 25.9% of IPv6 and 21.4% of IPv4 ASes permit source address spoofing due to lacking OSAV protections. Additionally, the UDP protocol's inherent vulnerability, responding to unverified source addresses, enables attackers to spoof targets and direct amplified traffic floods, with our experiments confirming negligible packet loss during measurements.

III. Discovering Vulnerable ASes

A. Detecting ASes without ISAV

Autonomous systems (ASes) that have not deployed inbound source address validation (ISAV) are vulnerable to reflection amplification attacks. The key to identifying such ASes is determining whether ISAV filters spoofed-source packets. In the absence of probes within the target AS, the challenge becomes determining whether spoofed-source packets are received by hosts within an AS where they should not usually arrive. In this work, we leverage the Internet Control Message Protocol (ICMP) [16] to detect such "receptions" and

employ a cross-verification approach to confirm whether these "receptions" are indeed invalid.

Given that we control two vantage points, we could implement a reflection test by one vantage point, which can spoof the source address as the other's, thus triggering hosts within vulnerable ASes to respond to the target. By observing the received traffic at the target, we can infer whether those hosts have "received" packets.

To quickly explore ISAV deployment across multiple ASes, we can incrementally set the TTLs of packets, triggering routers along the path from a vantage point to a targeted host to send *ICMP Time Exceeded* messages. Similarly, the host will send *ICMP Destination Unreachable* or *ICMP Echo Reply* packets according to UDP-based or ICMP-based implementations of *traceroute*. The main distinction here is that the vantage point falsifies the source IP address of the probing packets to that of the target. As a result, if the target receives an ICMP error message, it can be determined that the corresponding router or host "receives" the spoofed packets.

B. Validation with Traceroute

Simply detecting that routers or targeted hosts are "receiving" spoofed packets does not necessarily imply that the corresponding ASes have not deployed ISAV. This is because the traffic from both the vantage point and the target may traverse the same routers on their way to the hosts. From the perspective of the ASes where these routers are located, packets originating from both the vantage point and the target appear to be from the same source in terms of network topology. Under such circumstances, ISAV cannot effectively differentiate or filter the spoofed packets. Therefore, we must rely on *traceroute* data of both the vantage and target to identify truly invalid "receptions".

As illustrated in Fig. 1, this paper concentrates on following two positional relationships among an AS and routing paths of a vantage point and a target to a host.

- **Type 1**: As illustrated in Fig. 1 (a) and (b), all routers within AS_i are in the path from vantage point V to host H and the path from target T to host H. In the case of Fig. 1 (a), even if AS_i has deployed ISAV, it cannot distinguish spoofed packets that appear to originate from the same source as the target in the network topology. Furthermore, the results of routers in Fig. 1 (a) and Fig. 1 (b) "receiving" spoofed packets along the route from vantage point V to host H are indistinguishable. Therefore, we cannot definitively determine whether AS_i has deployed ISAV in such scenarios.
- **Type 2**: As shown in Fig. 1 (c) and (d), there is a router R_j within AS_i that is on the path from vantage point V to host H but not on the path from target T to host H. In this scenario, if R_j "receives" the spoofed packet, it can be concluded that AS_i has not deployed ISAV.

C. Identifying Practical Reflectors

Scanning Method. According to IPASN (October 2024) [21], there are currently 30,368 ASes enabled in the IPv6

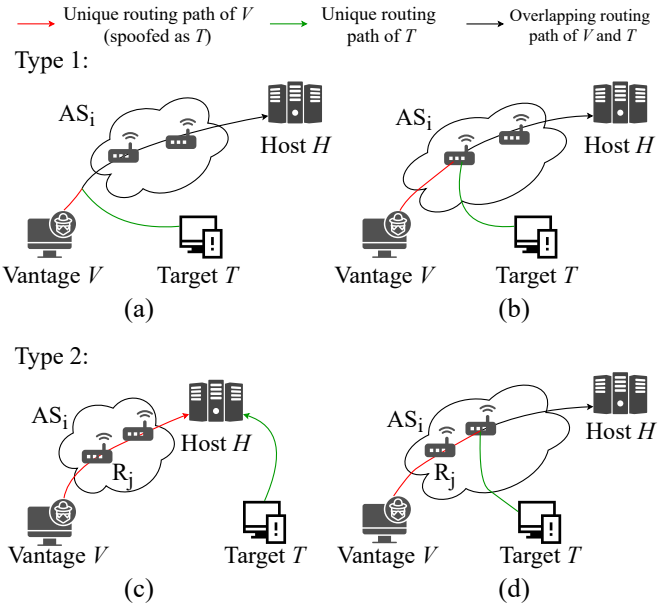

Fig. 1. Three types of positional relationships among AS AS_i and routing paths of vantage point V and target T to host H. The red, green, and black arrows represent the unique path from V to H with the source address spoofed as T, the unique path from T to H, and the overlapping path of the two, respectively.

network, covering 225,419 BGP prefixes. To achieve AS-level coverage, we refer to [17], performing prefix consolidation and random-byte generation for each remaining BGP prefix to collect IPv6 addresses as targeted hosts (cf. Alg. 1 Line 2). Similarly to [15], we utilize Yarrp [22] to randomly permute a targeted host × TTL space to trigger reflection at in-path routers in a stateless manner. The corresponding algorithm is presented in Alg. 1.

Algorithm 1: Discovering ASes without ISAV

```
   Input: BGP prefixes Prefix
   Output: Vulnerable ASes Vul_ASes
1  Initially Vul_ASes = [ ];
2  Targeted_hosts = Random_generation(Prefixes);
3  RT_vantage = Yarrp(Targeted_hosts, vantage,
     vantage) ;
4  RT_spoof = Yarrp(Targeted_hosts, vantage, target) ;
5  RT_target = Yarrp(Targeted_hosts, target, target) ;
6  for host ∈ Targeted_hosts do
7      AS_host = IP_to_AS(host) ;
8      if AS_host ∉ Vul_ASes then
9          Group_spoof = Filter(Routers_spoof, host) ;
10         Group_target = Filter(Router_target, host) ;
11         for router ∈ Group_spoof do
12             if router ∉ Group_target then
13                 Vul_ASes.append(AS_host) ;
14                 break ;
15             end
16         end
17     end
18 end
```

The vantage point first performs a standard *traceroute* to

identify ICMP-responsive routers (Line 3), then spoofs the target's address to detect routers "receiving" forged packets (Line 4). By comparing the spoofed and target's *traceroute* results (Lines 5-18), Type 2 vulnerable ASes are identified. Ideally, routers in ASes before the first ISAV-enabled AS will respond to spoofed packets, enabling multi-AS ISAV detection in one attempt. The vantage point's normal and spoofed paths help identify ISAV-deploying ASes (those responding normally but blocking spoofed traffic), though only the first such AS per host can be confirmed.

Notably, while routers respond with ICMP Time Exceeded messages, this doesn't guarantee UDP-based amplification capability, necessitating further validation for reflection amplification risks.

Identifying Vulnerable Servers. As discussed in Key Observation in Sec I, the third crucial condition for hosts to serve as amplifiers is the presence of vulnerable protocols on those hosts. This paper focuses on reflection amplification attacks facilitated by three common UDP services: DNS, NTP, and SNMP [1]. To determine whether these services are available on active addresses collected from IPv6 Hitlists [23] and proactive target generation scans [24], we use Xmap [25] to probe for service presence of DNS, NTP, and SNMP.

Reflection Amplification Tests. We have control over both a vantage point and a target, allowing the vantage point to perform source address spoofing. The vantage point can send service request packets with source addresses spoofed to that of the target, directed at service-providing hosts. Meanwhile, the target listens for the reflected traffic from the hosts. By analyzing this traffic, we can identify various amplifiers and match ASes containing them. This method not only reveals the presence of amplifiers but also enhances our understanding of the vulnerabilities inherent in certain ASes. By systematically probing these services and analyzing the responses, we can effectively map the landscape of potential threats and inform mitigation strategies for vulnerable networks.

IV. Measurement Results

A. Data Collection

In October 2024, we conducted five rounds of comprehensive AS-level ISAV deployment detection, using two controlled vantage points in China and Canada. For ASes found lacking ISAV, we collected active addresses followed by service probing and reflection amplification verification. To minimize network congestion, we scanned vulnerable ASes at a rate of 10 Kps for detecting "receives" and 1 Kps for cross-validation, with the initial TTL set to 4 to avoid triggering frequent ICMP reports from routers near the vantage point. Meanwhile, service probing was conducted at 10 Kps. To reduce interference with the target and nearby routers, we conducted reflection amplification validation at a rate of only 0.1 Kps. After each scan, we summarized the target response status and the corresponding AS distribution.

B. Are There Actually Amplifiers in ASes without ISAV?

The Deployment of ISAV. After conducting five rounds of AS-level scanning, we measured 7,269 ASes. Among them, routers in 5,914 ASes are found to receive spoofed packets. By correlating these results with the target's *traceroute* data, we confirm that 4,460 ASes (61.36%) are classified as not deploying ISAV. Meanwhile, 1,560 ASes (21.46%) are found to have deployed ISAV, and the status of the remaining 1,249 ASes (17.18%) remains uncertain.

Compared to earlier studies, [13] and [12] found that 68.6% of the tested ASes either lacked or only partially deployed ISAV in 2019. Additionally, [26] reported that 69.8% of tested ASes were without ISAV in 2021. Similarly, in 2022, [15] demonstrated that ISAV deployment was either inadequate or only partially completed in 67.37% of tested ASes, with full deployment occurring in 20.87%. Our findings are consistent with the results of these studies.

Amplifiers. We collected approximately 47 million active addresses within these ASes lacking ISAV. Tab. I presents the service probing and reflection amplification validation results for these addresses. Among the 4,460 ASes lacking ISAV, we successfully identified exploitable reflectors in 3,507 ASes (78.63%), providing strong evidence to support our hypothesis that "ASes lacking ISAV are likely to harbor amplifiers."

Vulnerable ASes Distribution. As shown in Fig. 2, with data from [27], we find that the 4,460 ASes without ISAV are distributed across 168 countries. Among them, detection rates for vulnerable ASes in countries and regions like the Middle East and Central Africa are relatively low. This may be because most of these countries and areas have meager IPv6 adoption rates—less than 1% [4].

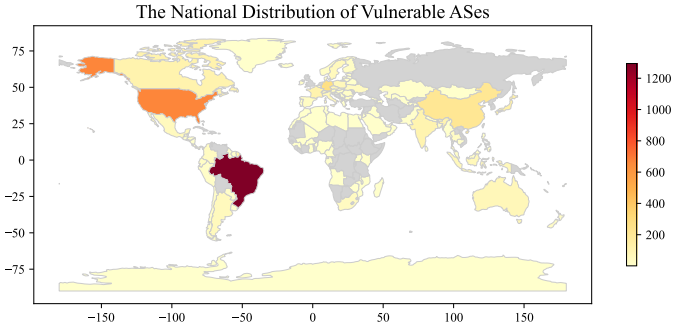

Fig. 2. Heatmap of the distribution of vulnerable ASes around the world.

It can be concluded that the scope of reflection amplification attacks in IPv6 is relatively wide and warrants caution.

C. How Harmful Amplifiers are in IPv6?

The key metric for measuring reflection amplification attacks is the amplification factor. In calculating the amplification factor, we assume a single amplifier, disregard the time used for query and response timing, and include Ethernet headers when determining packet size. The bandwidth amplification factor (BAF) is thus calculated as follows:

$$BAF = PAF \cdot SAF = \frac{\#\ of\ response\ packets}{\#\ of\ query\ packets} \cdot \frac{Response\ packet\ size}{Query\ packet\ size}. \quad (2)$$

TABLE I
THE NUMBER OF OPEN SERVERS, VERIFIED AMPLIFIERS, SUCCESSFUL EXPLOITATIONS, AND ASES COVERED FOR EACH PROTOCOL.

Protocols	# of Hosts Providing Services	# of Amplifiers	Successful Exploitations (%)	# of ASes Covered
DNS	85,952	81,833	95.2	596
NTP	34,022	21,691	63.8	1,555
SNMPv1	3,274	2,803	85.6	539
SNMPv2	3,499	3,016	86.2	552
SNMPv3	53,935	49,952	92.6	3,029
Total	158,815	141,006	88.8	3,507

We only select some effective amplifiers to present the amplification performance to minimize the impact on real-world networks. The measure results and test scenarios for DNS, NTP, and SNMP are summarized in Tab. II.

DNS: In IPv4, a standard method for conducting reflection amplification using DNS is through name lookups (e.g., A or MX records). With the DNS extension (EDNS0) allowing UDP response packets larger than 512 bytes, the amplification factor generally ranges from 28 to 54 [1]. In our experiment, we test four types of queries: 15, 16, 28, and 255, corresponding to "MX", "TXT", "AAAA", and "ANY" query types, respectively. When DNSSEC and EDNS0 are enabled with "udp payload size" set to 65535 to break the packet size limit, we achieve an average amplification factor of 32.51 and a maximum of 42.87 for querying Microsoft's "ANY" records.

NTP: Similar to IPv4, vantages in IPv6 can exploit the *monlist* request supported by NTP servers, which returns up to 100 packets with data on recent clients. In our experiment, we do not attempt to maximize the amplification factor for each NTP server to minimize disruption to actual network devices. Instead, we measure the network's existing amplification factors of NTP servers and test an average BAF of 665.38. In most cases, a *monlist* request can increase request traffic by up to $717.14\times$ while responding with no more than 100 packets. This aligns with the explanation regarding *monlist*. Surprisingly, some servers can amplify by $4267\times$ with 595 packets responded, which is the highest BAF measured.

SNMP: SNMP has three versions, with SNMPv2 and SNMPv3 supporting the *GetBulk* command for large responses. For SNMPv1, a crafted query packet can achieve an average amplification factor of 27.03 and a maximum of 43.18. SNMPv2, with higher *max_repetitions*, reaches an average of 115.26 and a peak of 481.27. SNMPv3, however, includes authentication, limiting its amplification factor to a maximum of 20.07 and an average of 4.46. As SNMPv3 introduces an authentication mechanism, *GetBulk* is valid for vantages, and the maximum amplification factor for SNMPv3 is limited to 20.07, with an average of only 4.46. Additionally, SNMPv3 exposes router vendor information in error messages [28].

According to Tab. I, DNS has the highest exploitation success rate but is concentrated in specific areas. SNMPv3 is widely deployed but has low amplification, while SNMPv2 has higher amplification but less exposure. NTP has a lower exploitation success rate but offers the highest amplification and wider distribution, making it the most ideal amplifier.

V. DISCUSSION

IP Deactivation. IPv6 addresses, especially on mobile devices, frequently change due to network reconnections, leading to short lifespans. Our one-week liveness test found only 4.71% of previously active addresses remained reachable, confirming rapid turnover. However, this volatility does not affect AS-level vulnerability detection, as reassigned addresses stay within the same AS. For reflection attacks, IPv6's dynamic nature complicates attacker reconnaissance (requiring repeated scanning) but also enhances anonymity, hindering victim forensics.

Countermeasures. Reflection attacks rely on source IP spoofing. Outbound (OSAV) and inbound (ISAV) source address validation can block spoofed packets: OSAV prevents local-source spoofing, while ISAV filters illegitimate inbound traffic. Despite their importance, many ASes lack full deployment, especially for IPv6. Promoting ISP adoption is crucial, though ISAV only blocks cross-network attacks—internal monitoring remains essential. What's more, hosts can reduce exposure by restricting access (e.g., DNS servers for internal users only), using authentication (e.g., SNMP passwords), or rate-limiting responses (e.g., DNS RRL). Passive defenses like ACL filtering and DPI can block malicious traffic by IP or content patterns, disrupting attack pathways.

Limitations. We acknowledge certain limitations in this study. When detecting ASes lacking ISAV, we infer whether routers accept spoofed packets based on ICMP timeout messages received by the target. However, not all routers reliably return ICMP timeout messages. Many gateways are configured by default not to return ICMP messages for security reasons. Consequently, monitoring from the target's side alone may miss routers receiving spoofed packets.

Additionally, this paper relies on only one probe capable of source address spoofing and one passive listening to obtain all detection results, resulting in limited coverage. Once forged packets are blocked by an AS with ISAV deployed, assessing ISAV deployment in subsequent ASes becomes challenging.

Ethical Consideration. We maintain **anonymity** by only disclosing country-level vulnerable ASes, following established practices. Our scanning uses **low-impact methods**: standard ICMP Echo requests, distributed traceroute (starting at TTL=4 with randomized targets), and minimal service checks to avoid network disruption. We're coordinating with local network administrators to **responsibly disclose** identified IPv6 reflection vulnerabilities given their broad implications.

TABLE II
BANDWIDTH AMPLIFICATION FACTORS FOR EACH PROTOCOL.

Protocols	Query Packet Size (B)	Response Packet Size (B)		# of Response Packets		BAF	
		Avg	Max	Avg	Max	Avg	Max
DNS	100	1083.69	1429	3	3	32.51	42.87
NTP	70	501.15	502	92.9	595	665.38	4267
SNMPv1	96	1297.81	1396	2.02	3	27.31	43.63
SNMPv2	97	1318.45	1505.9	8.48	31	115.26	481.27
SNMPv3	142	211.03	190	3	15	4.46	20.07

VI. CONCLUSION

This paper investigates reflection amplification attacks in IPv6 networks. Using only two probes, we have efficiently assessed the vulnerabilities of ASes and successfully identified reflection amplifiers within these ASes. Our findings confirm the widespread presence of reflection amplifiers in the current IPv6 networks. Additionally, we have analyzed the abuse potential of three UDP-based network protocols in reflection amplification attacks, highlighting this "grey rhino" problem.

ACKNOWLEDGMENT

This work is supported by the National Natural Science Foundation of China No.62472434 and China Postdoctoral Science Foundation No.2023TQ0089.

REFERENCES

[1] C. Rossow, "Amplification hell: Revisiting network protocols for ddos abuse." in *In the proceeding of the 30th USENIX Security Symposium (USENIX Security 14)*, 2014, pp. 1–15.

[2] S. Kottler, "February 28th ddos incident report," 2018. [Online]. Available: https://github.blog/2018-03-01-ddos-incident-report/

[3] A. C. Security, "How to deal with ddos attacks on a global scale," 2019. [Online]. Available: https://www.alibabacloud.com/blog/how-to-deal-with-ddos-attacks-on-a-global-scale_595641

[4] Akamai, "Ipv6 adoption visualization," 2024. [Online]. Available: https://www.akamai.com/internet-station/cyber-attacks/state-of-the-internet-report/ipv6-adoption-visualization

[5] X. Li, D. Wu, H. Duan, and Q. Li, "Dnsbomb: A new practical-and-powerful pulsing dos attack exploiting dns queries-and-responses," in *Proceedings of the 2024 IEEE Symposium on Security and Privacy (SP)*. IEEE Computer Society, 2024, pp. 253–253.

[6] W. Xu, X. Li, C. Lu, B. Liu, H. Duan, J. Zhang, J. Chen, and T. Wan, "Tsuking: Coordinating dns resolvers and queries into potent dos amplifiers," in *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security*, 2023, pp. 311–325.

[7] G. C. Moura, S. Castro, J. Heidemann, and W. Hardaker, "Tsuname: exploiting misconfiguration and vulnerability to ddos dns," in *Proceedings of the 21st ACM Internet Measurement Conference*, 2021, pp. 398–418.

[8] C. Deccio, A. Hilton, M. Briggs, T. Avery, and R. Richardson, "Behind closed doors: a network tale of spoofing, intrusion, and false dns security," in *Proceedings of the ACM Internet Measurement Conference*. ACM, 2020, pp. 65–77.

[9] 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 *Proceedings of the Passive and Active Measurement: 21st International Conference*. Springer, 2020, pp. 107–121.

[10] J. Czyz, M. Kallitsis, M. Gharaibeh, C. Papadopoulos, M. Bailey, and M. Karir, "Taming the 800 pound gorilla: The rise and decline of ntp ddos attacks," in *Proceedings of the 2014 Conference on Internet Measurement Conference*, 2014, pp. 435–448.

[11] CAIDA, "Ip spoofing," 2024. [Online]. Available: https://spoofer.caida.org/summary.php

[12] 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 *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, 2019, pp. 465–480.

[13] CAIDA, "Spoofer project," 2022. [Online]. Available: https://www.caida.org/projects/spoofer

[14] R. Beverly, A. Berger, Y. Hyun, and K. Claffy, "Understanding the efficacy of deployed internet source address validation filtering," in *Proceedings of the 9th ACM SIGCOMM conference on Internet measurement*, 2009, pp. 356–369.

[15] L. Pan, J. Yang, L. He, Z. Wang, L. Nie, G. Song, and Y. Liu, "Your Router is My Prober: Measuring IPv6 Networks via ICMP Rate Limiting Side Channels," in *Proceedings of the 30th Annual Network and Distributed System Security Symposium, NDSS 2023*, no. March, 2023.

[16] M. Gupta and A. Conta, "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification," RFC 4443, Mar. 2006. [Online]. Available: https://www.rfc-editor.org/info/rfc4443

[17] G. Song, J. Yang, Z. Wang, L. He, J. Lin, L. Pan, C. Duan, and X. Quan, "Det: Enabling efficient probing of ipv6 active addresses," *IEEE/ACM Transactions on Networking*, vol. 30, no. 4, pp. 1629–1643, 2022.

[18] K. Bock, A. Alaraj, Y. Fax, K. Hurley, E. Wustrow, and D. Levin, "Weaponizing middleboxes for {TCP} reflected amplification," in *Proceedings of the 30th USENIX Security Symposium (USENIX Security 21)*, 2021, pp. 3345–3361.

[19] Q. Lone, M. Luckie, M. Korczyński, and M. Van Eeten, "Using loops observed in traceroute to infer the ability to spoof," in *Passive and Active Measurement: 18th International Conference, PAM 2017, Sydney, NSW, Australia, March 30-31, 2017, Proceedings 18*. Springer, 2017, pp. 229–241.

[20] L. Müller, M. Luckie, B. Huffaker, K. Claffy, and M. Barcellos, "Challenges in inferring spoofed traffic at ixps," in *Proceedings of the 15th International Conference on Emerging Networking Experiments and Technologies*, 2019, pp. 96–109.

[21] H. Asghari and A. Noroozian, "Ip address to autonomous system number lookups," 2024. [Online]. Available: https://pypi.org/project/pyasn/

[22] R. Beverly, "Yarrp'ing the internet: Randomized high-speed active topology discovery," in *Proceedings of the 2016 Internet Measurement Conference*, 2016, pp. 413–420.

[23] B. Hou, Z. Cai, K. Wu, T. Yang, and T. Zhou, "Search in the expanse: Towards active and global ipv6 hitlists," in *IEEE INFOCOM 2023-IEEE Conference on Computer Communications*. IEEE, 2023, pp. 1–10.

[24] T. Yang, Z. Cai, B. Hou, and T. Zhou, "6forest: an ensemble learning-based approach to target generation for internet-wide ipv6 scanning," in *IEEE INFOCOM 2022-IEEE Conference on Computer Communications*. IEEE, 2022, pp. 1679–1688.

[25] X. Li, B. Liu, X. Zheng, H. Duan, Q. Li, and Y. Huang, "Fast ipv6 network periphery discovery and security implications," in *2021 51st Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. IEEE, 2021, pp. 88–100.

[26] T. Dai and H. Shulman, "Smap: Internet-wide scanning for spoofing," in *Proceedings of the 37th Annual Computer Security Applications Conference*, 2021, pp. 1039–1050.

[27] Maxmind, "Geolite2 free geolocation data," 2024. [Online]. Available: https://dev.maxmind.com/geoip/geolite2-free-geolocation-data/

[28] T. Albakour, O. Gasser, R. Beverly, and G. Smaragdakis, "Third time's not a charm: Exploiting SNMPv3 for router fingerprinting," *Proceedings of the ACM Internet Measurement Conference, IMC*, pp. 150–164, 2021.