Poster: Quantifying the Proportion of Hijacked Prefixes Among the Identified Prefix Hijackers

Ebrima Jaw[†], Moritz Müller[†*], Cristian Hesselman[†*], Lambert Nieuwenhuis[†]
[†]University of Twente, Enschede, The Netherlands
[*]SIDN Labs, Arnhem, The Netherlands

***Abstract*—The Border Gateway Protocol (BGP) remains susceptible to prefix hijacks due to its trust-based nature and lack of default robust authentication mechanisms. Prefix hijacks are unintentional or malicious announcements of prefixes allocated to other ASes. Although prefix hijacks are primarily associated with misconfigurations, they remain a significant security threat. For instance, the recent hijacking and route leak incident involving Cloudflare made their DNS resolver unreachable for some networks for about 8 hours. Some ASes perform hijacks frequently and for longer duration. We revisited these "serial hijackers" in 2024 and validated some of the potential serial hijackers with external data. However, neither the original study from 2019 nor ours dug deeper to understand the impact and goal of serial hijackers. This study fills this gap and shows that 22.9% of the announcements were RPKI-*invalid*, raising new questions about the intent of the hijack. Finally, we show that these invalid announcements still reach many networks on the Internet, demonstrating that many ASes are not doing RPKI route origin validation, thereby compromising the Internet's stability and security.**

I. Introduction

Autonomous networks on the Internet use the Border Gateway Protocol to exchange reachability information, such as IP prefixes (172.16.0.0/16) and path attributes (e.g., origin and next-hop) with their peers [1]. BGP's decentralized and trust-based nature has facilitated the Internet's scalability and linear growth. However, these attributes make BGP vulnerable to misconfigurations and prefix hijacks, often exploited by malicious actors, which remain a security threat to the Internet.

Bogus routes from accidental or malicious prefix hijacks can propagate across the Internet without IRR and RPKI route filtering, leading to the interception of sensitive data, service downtime, and interruption of critical network traffic. For example, the recent hijacking and route leak incident involving a more specific prefix (1.1.1.1/32) of Cloudflare's 1.1.1.0/24 made the DNS resolver unreachable for over 300 networks in 70 countries for about 8 hours before fully resolved [2].

Irrespective of the various available route filtering, monitoring, and detection methods [3], one-time prefix hijacks happen frequently. However, in 2019, Testart et al. defined the concept of *serial hijackers* as networks that repeatedly hijacked prefixes for longer durations [3], which we recently reproduced and extended to investigate the current dynamics of serial hijacking activities on the Internet [4]. Similar to the original study, we found that 766 networks show the behavior of serial hijacking (*flagged ASes*), and we validated some of them using supplementary data sources. However, neither the original study nor ours characterize these potential serial hijackers further and try to understand their motivation. This work presents the first steps to bridge this gap. More concretely, this study presents an automated pipeline that (i) quantifies the proportion of hijacked prefixes among our flagged ASes, (ii) identifies the victim networks and shows the most vulnerable resources to prefix hijacks, and (iii) investigates potential legitimate and malicious hijacking activities.

978-3-903176-64-5

II. Data sets and Methodology

In addition to the *RIR delegation files*, *longitudinal MOAS*, and *BGP dataset* discussed in [4], we have also used RIPENCC's RPKI snapshots between January 2019 and November 2023 to determine the route origin validation (ROV) status of hijacked prefixes [5]. Our automated data collection and analysis pipeline enables us to extract all the relevant BGP data for our *flagged ASes*, such as the *announcement dates*, *originated prefixes*, *origin ASes*, and *peer counts*. Then, we used the prefix origin pairs of all the announcements we observed of our flagged ASes to determine their RPKI status between 2019-2023. Next, we filtered out the unique RPKI-*invalid* announcements by the flagged ASes. To identify the *victim networks*, we used prefix to AS mapping to look for the ASes assigned with the invalid prefixes. Finally, since current literature shows that providers can announce their customers' prefixes for visibility [6], we augment CAIDA's ASRank data with the AS relationship dataset to identify potential P2C or C2P relationships among the flagged and victim ASes.

III. Preliminary Results

Our flagged ASes announced 92.7K unique prefixes, and 19K of these announcements were RPKI-*invalid*, 1.1K of which were MOAS-*invalid* announcements [6]. Table I shows the five-year aggregated announcements among our flagged ASes. We observed 42.8% *unknown* announcements, indicating that the originating ASes have no registered ROAs for these prefixes. Interestingly, 22.9% of the announcements were *invalid*, which signifies existing ROAs for these announcements but for different ASes than our flagged ASes.

Next, we compute the average ($\overline{X}$), weighted average ($\overline{X}_w$)[1], median ($\tilde{X}$), and standard deviation (σ) of peer counts (pcs) for the RPKI statuses to quantify the propagation of

[1]Divide the total pcs of each RPKI status by the sum of all pcs to account for the relative contribution of each RPKI status to the total pcs distribution.

these announcements on the Internet. The descriptive statistical

TABLE I: Aggregate announcements among flagged ASes

Year	Invalid ASN	Invalid Length	Unknown	Valid
2019	4,259 (32.6%)	2,084 (15.9%)	6,033 (46.1%)	699 (5.3%)
2020	3,442 (35.6%)	1,576 (16.3%)	3,620 (37.4%)	1,032 (10.7%)
2021	6,309 (19.7%)	1,629 (5.1%)	18,279 (57.2%)	5,740 (18.0%)
2022	2,194 (16.8%)	1,901 (14.5%)	6,607 (50.5%)	22,390 (18.3%)
2023	2,795 (9.8%)	3,665 (12.9%)	6,518 (22.9%)	15,447 (54.3%)
Average	3,800 (22.9%)	2,171 (12.9%)	88,211 (42.8%)	55,062 (21.3%)
Descriptive Statistics: Weighted avg. median and standard deviation of peer counts among the RPKI statuses				
$\overline{X}$ and ($\overline{X}_w$%)	441 (24.1%)	124 (3.9%)	260 (30.7%)	567 (41.2%)
Median ($\tilde{X}$)	470.5	313.0	430.0	628.0
Std (σ)	98	110	104	99

summary at the bottom of Table I provides a more nuanced understanding of the propagation of announcement types. It shows that *valid announcements* among our flagged ASes are predominantly propagated among peers, which is desirable and expected for the security and stability of the Internet. However, we observed a substantial weighted average peer count of 30.7% and 24.1% for *unknown* and *invalid* announcements. Additionally, despite the variability in the values of the *unknown* and *invalid* announcements, which could be due to partial RPKI route validation on the Internet, they are somewhat accepted among peers. However, the variability in σ, especially for *invalid length*, underscores that the *invalid length* announcements are less propagated among peers and are diversely handled across the Internet. Lastly, although the numbers presented in Table I might seem unsurprising considering the low adoption rate of RPKI, they still pose potential security threats to the Internet infrastructure, as discussed in the use case presented at the end of this section.

TABLE II: Invalid prefix announcements to AS mapping

RIPENCC	ARIN	APNIC	AFRINIC	LACNIC
7,984 (44.4%)	5,559 (30.9%)	2,038 (11.3%)	1,455 (8.1%))	933 (5.2%)
Announcement share of *unallocated prefixes*				
1,175 (14.7%)	2,377 (42.8%)	296 (14.5%)	123 (8.45%)	284 (30.4%)

Our *invalid prefix to AS mapping* results reveal that these potentially hijacked prefixes belong to 2221 unique ASes (victims) from all the five RIRs, and Table II shows five years of aggregated results. Notely, 7.984K (44.4%) and ≈5.6K (30.9%) of these invalid announced prefixes were allocated to ASes within RIPENCC and ARIN region, respectively. AFRINIC and LACNIC have the lowest at 8.1% and 5.2%. Additionally, we found that 4.3K of the invalid prefixes had no allocation details from our RIR delegation files. About 42.9%, 30.4%, and 14.7% of these potentially hijacked *unallocated prefixes* were delegated to the ARIN, LACNIC, and RIPENCC. Irrespective of the many prefixes managed by RIPENCC and ARIN, these correlated results show that their resources are more vulnerable to hijacking incidents.

Interestingly, we found that most of these invalid announced prefixes are more specific than the allocated prefixes from the RIRs. For instance, about 13K (82.3%) and 3K (90.7%) of the invalids had prefix lengths of /24 and /32-47 for IPv4 and IPv6, respectively. Figure 1 summarized our results.

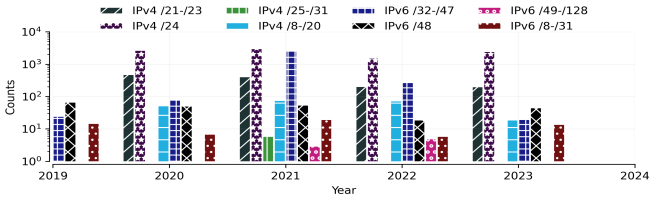

Fig. 1: Distributions of announced invalid prefix lengths

a) Use Case: Invalid Announcements: Figure 2 shows that AS33696 (*a flagged AS*) has a daily average of 82.49% invalid unique prefix announcements, indicating it has no registered ROAs for these prefixes. In contrast, its valid announcements (6.74%) only started from 2022 to Nov. 2023. Suprisingly, 81.20% of its unique announced prefixes belong to 64 different ASes in 5 regions, such as ripencc (53.7%), arin (29.62%), afrinic (12.04%), lacnic (3.70%), apnic (0.93%). Although these behaviors are suspicious, we could not determine the intent. However, AS33696 is among the 135 of our flagged ASes that went out of operation between 2019 to 2024.

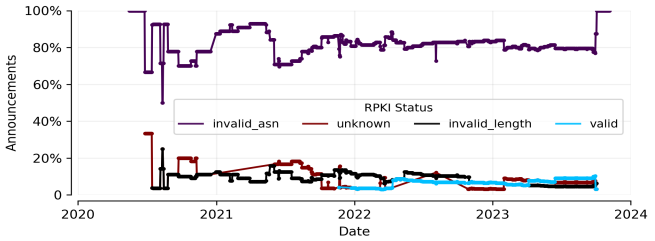

Fig. 2: RPKI status of AS33696 announcements over time.

b) Lessons Learned so far: Determining the root causes of prefix hijacking among our flagged ASes is challenging due to the dynamic and complex implicit relationships among ASes. We also learned that IP leasing companies could be exacerbating the issues with their ROAs handling methods. Finally, as in previous studies, our analysis shows again that ROV is deployed incompletely, which leaves significant networks unprotected.

IV. Conclusion and Work in Progress

This study attempted to better understand the behavior of serial hijackers. We showed that even RPKI-*invalid announcements* are seen by many route collector peers, indicating that ROV has yet to be deployed sufficiently. Furthermore, the identified RPKI-*invalid* announcements among serial hijackers remain a significant security vulnerability. Therefore, we plan to further study these *invalid* announcements to determine malicious and potential legitimate use cases to reduce false positives, thereby improving current hijacking detection methods. Finally, we have found indications of IP leasing during our analysis, which could lead to falsely classifying an event as a hijack. Thus, we intend to quantify its effect on invalid or unknown announcements and provide BCPs for handling ROAs for IP leasing companies.

REFERENCES

[1] B. Al-Musawi, P. Branch, and G. Armitage, "BGP Anomaly Detection Techniques: A Survey," *IEEE Communications Surveys & Tutorials*, vol. 19, no. 1, pp. 377–396, 2017.

[2] B. Herdes, M. Zhang, and T. Ryan. Cloudflare 1.1.1.1 incident on june 27, 2024. [Online]. Available: https://blog.cloudflare.com/cloudflare-1111-incident-on-june-27-2024

[3] C. Testart, P. Richter, A. King, A. Dainotti, and D. Clark, "Profiling BGP serial hijackers: Capturing persistent misbehavior in the global routing table," in *Proceedings of the Internet Measurement Conference*, ser. IMC '19. New York, NY, USA: Association for Computing Machinery, 2019-10-21, pp. 420–434.

[4] E. Jaw, M. Müller, C. Hesselman, and L. Nieuwenhuis, "Serial BGP hijackers: A reproducibility study and assessment of current dynamics," in *2024 8th Network Traffic Measurement and Analysis Conference (TMA)*, 2024-05, pp. 1–10.

[5] RIPE NCC. Rpki snapshots/. [Online]. Available: https://ftp.ripe.net/rpki/

[6] K. Z. Sediqi, A. Feldmann, and O. Gasser, "Live long and prosper:analyzing long-lived MOAS prefixes in BGP," 2023-07-17. [Online]. Available: http://arxiv.org/abs/2307.08490