Understanding Route Origin Validation (ROV) Deployment in the Real World and Why MANRS Action 1 Is Not Followed

Lancheng Qin*‡, Li Chen†, Dan Li*†, Honglin Ye*, and Yutian Wang*
*Tsinghua University,†Zhongguancun Laboratory,‡BNRist
{qlc19, yehl20, wangyt21}@mails.tsinghua.edu.cn, lichen@zgclab.edu.cn, tolidan@tsinghua.edu.cn

***Abstract*—BGP hijacking is one of the most important threats to routing security. To improve the reliability and availability of inter-domain routing, a lot of work has been done to defend against BGP hijacking, and Route Origin Validation (ROV) has become the best current practice. However, although the Mutually Agreed Norms for Routing Security (MANRS) has been encouraging network operators to at least validate announcements of their customers, recent research indicates that a large number of networks still do not fully deploy ROV or propagate illegitimate announcements of their customers. To understand ROV deployment in the real world and why network operators are not following the action proposed by MANRS, we make a long-term measurement for ROV deployment and further find that many non-compliant networks may deploy ROV only at part of customer interfaces, or at provider or peer interfaces. Then, we present the first notification experiment to investigate the impact of notifications on ROV remediation. However, our analysis indicates that none of the notification treatments has a significant effect. After that, we conduct a survey among network operators and find that economical and technical problems are the two major classes of reasons for non-compliance. Seeking a realistic ROV deployment strategy, we perform large-scale simulations, and, to our surprise, find that not following MANRS Action 1 can lead to better defence of prefix hijacking. Finally, with all our findings, we provide practical recommendations and outline future directions to help promote ROV deployment.**

I. Introduction

Nowadays, there are more than 70,000 Autonomous Systems (ASes) in the Internet. ASes use Border Gateway Protocol (BGP) to exchange routing information and determine the best route to each destination prefix. However, due to the lack of built-in security mechanism, BGP is vulnerable to various malicious attacks, among which BGP hijacking is the most common and harmful. A BGP hijacking, or BGP prefix hijacking, is when an attacker tries to hijack the traffic flowing to another network by announcing prefixes belonging to this network in BGP. It is mostly used for traffic disruption [1], DDoS attack [2], eavesdropping attack [3], sending spam [4], [5], or stealing crypto currencies [6], [7].

To prevent BGP hijacking and enhance the reliability of inter-domain routing, many BGP security mechanisms [8], [9], [10], [11], [12], [13], [14] have been proposed from both academia and industry. However, few of them are widely deployed in practice. As a result, BGP hijacking still remains a significant threat to today's Internet. In 2021, BGPStream reported 775 BGP hijacking incidents, some of which caused significant disruptions around the world [15]. For example, on October 25, 2021, AS 212046 was reported to hijack 3,786 prefixes, creating conflicts with 972 ASes in 42 countries [16].

To promote the deployment of BGP security mechanisms, the Mutually Agreed Norms for Routing Security (MANRS) initiative [17], supported by Internet Society (ISOC), provides practical fixes to mitigate BGP hijacking and proposes four actions for network operators [18]. Specifically, MANRS Action 1 requires that network operators must check whether the BGP announcements received from their customers are correct by using Internet Routing Registries (IRRs) [19], Resource Public Key Infrastructure (RPKI) [13], or other manual route filters. Both IRRs and RPKI maintain the objects that document which ASes are allowed to announce which IP address spaces. However, most IRRs do not have strict authentication mechanisms to prevent the users from entering fake data. In contrast, RPKI cryptographically validates its objects (*i.e.*, Route Origin Authorizations (ROAs)) to guarantee the authenticity of its data. Therefore, Route Origin Validation (ROV) [20], the mechanism specifically designed to identify BGP hijacking using ROAs, has become the best current practice to secure BGP in recent years.

However, according to recent research on the measurement of ROV deployment [21], [22], [23], [24], more than 60% of ASes under test only selectively perform RPKI-invalid filtering depending on the interfaces [1] at which the BGP announcements arrive, or do not deploy ROV at all. Moreover, we notice that many ASes, including some Tier-1 ASes that claim to have already deployed ROV, may adopt ROV deployment strategies that are not compliant to MANRS Action 1. For example, AT&T (AS 7018) was reported to deploy ROV and discard RPKI-invalid prefixes only at peer interfaces, but not at customer interfaces [26]. But little is known about why network operators are not following MANRS Action 1.

To understand ROV deployment strategies used in practice and if practical ROV deployment is following MANRS

Network and Distributed System Security (NDSS) Symposium 2024
26 February - 1 March 2024, San Diego, CA, USA
ISBN 1-891562-93-2
https://dx.doi.org/10.14722/ndss.2024.24214
www.ndss-symposium.org

[1]The interface in this paper refers to a logical AS-level interface. For an AS, its multiple physical interfaces connected to the same AS are considered one interface. The provider/customer/peer interface refers to the interface connected to a provider/customer/peer [25].

Action 1, we first make an Internet-scale measurement to identify which ASes are accepting RPKI-invalid prefixes from customers, providers, or peers. Then, to the best of our knowledge, we conduct the first notification experiment to analyze the impact of notification on ROV remediation. To our disappointment, our results show that MANRS Action 1 is not followed by many ASes and notification cannot significantly improve their remediation rates of ROV. To identify the reasons for non-compliance, we conduct interviews and surveys among network operators in the Internet. From their replies, we summarize two major classes of reasons: economical and technical.

To promote ROV deployment and improve MANRS Action 1, we perform extensive simulations to determine the best deployment strategy. By following our proposed deployment strategy instead of MANRS Action 1, BGP prefix hijacking can be better prevented and some economical problems can be avoided. We also provide recommendations for backup and purchasing, and outline future directions for tackling technical challenges.

The main contributions of this work are the following:

- We perform a long-term measurement and identify 1,012 ASes (including 117 stub ASes and 895 non-stub ASes) that have passed RPKI-invalid prefixes. We find that 61.3% of these non-stub ASes do not deploy ROV at all customer interfaces and thus do not follow MANRS Action 1. Instead, 29.0% of the 1,012 ASes may deploy ROV at all provider interfaces, and 31.6% of them may deploy ROV at all peer interfaces.
- We present the first notification experiment to evaluate the impact of different notification treatments (including nudges and native language) on ROV remediation. However, our survival analysis indicates that none of the notification treatments can significantly improve the remediation rate of ROV.
- We conduct interviews with 5 network operators and find that MANRS Action 1 conflicts with their business interests. We then conduct a survey among more network operators and summarize that non-compliant networks are mainly due to the lack of time and effort, business conflicts, limited router capability, high operational overhead, technical bugs in ROV implementation, and technical limitations of ROV mechanism.
- We perform large-scale simulations to provide recommendations that help address the problems and improve MANRS Action 1. In particular, our simulations show the surprising result, that ROV at provider interfaces can work *better* in preventing the propagation of RPKI-invalid prefixes than ROV at customer or peer interfaces, without harming the business interests of transit networks. Therefore, we recommend that network operators first deploy ROV at provider interfaces, contrary to MANRS Action 1.

II. Background and Motivation

In this section, we first introduce the background of BGP hijacking, RPKI, and MANRS. We then propose our motivation and methodology of this work.

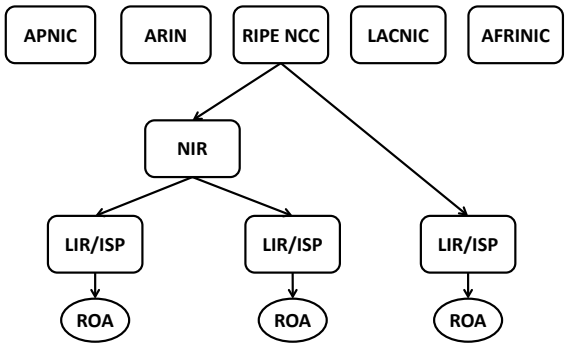

Figure 1: The structure of Resource Public Key Infrastructure.

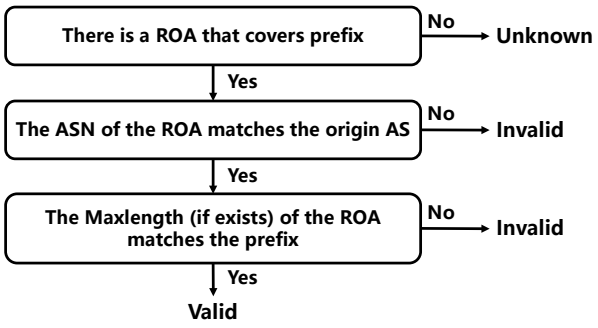

Figure 2: Validation process of Route Origin Validation.

A. BGP Hijacking

BGP hijacking is when an AS alters or forges any part of a BGP announcement to hijack the traffic forwarded to the victim. It can be caused by misconfiguration or malicious attacks. Specifically, there are two kinds of BGP hijacking, *i.e.*, prefix origin hijacking and path hijacking. Prefix origin hijacking means an AS maliciously announces the prefix of the victim in its BGP announcements. The path hijacking means an AS maliciously alter the AS path in its BGP announcements to place itself between the victim and other ASes. Compared with path hijacking, prefix origin hijacking is more commonly observed in the Internet [8]. It usually does not last long, but may pollute 90% of the Internet in less than two minutes [14]. Therefore, there has been continuous effort to mitigate prefix origin hijacking in both industry and academia.

B. Resource Public Key Infrastructure

To prevent prefix origin hijacking and strengthen routing security, various BGP security mechanisms have been proposed and RPKI has become the best current practice. As shown in Figure 1, RPKI is a hierarchical Public Key Infrastructure (PKI) to authorize the IP address space that can be legitimately announced by specific ASes via certificates. Since IP addresses and Autonomous Systems Numbers (ASNs) are allocated by five Regional Internet Registries (RIRs, *i.e.*, AFRINIC, APNIC, ARIN, LACNIC, and RIPE NCC), each RIR maintains and operates its own RPKI trust anchor with a root certificate. By using the root certificate, the RIR can

generate a signed certification for a Local Internet Registry (LIR, *e.g.*, network operator) or a National Internet Registry (NIR) with the resources (IP addresses and ASNs) assigned to the LIR or NIR. Eventually, network operators can sign ROA records to bind its own IP addresses with corresponding ASNs.

ROV is the application of RPKI to validate the authenticity for BGP announcements by using ROAs. When a router receives a BGP announcement, it can use ROV to determine whether there is a prefix origin hijacking by checking the prefix and the origin ASN of the BGP announcement against ROA records in RPKI repositories. Figure 2 illustrates the validation process of ROV: if there is no ROA that covers the prefix, the validation result is "unknown"; if there is a ROA that covers the prefix, but the origin ASN or the max-length is not matched, the validation result is "invalid"; if there is a ROA that matches both the prefix and the origin ASN, the validation result is "valid". If routers perform RPKI-invalid filtering, they will accept RPKI-valid and RPKI-unknown BGP announcements but discard RPKI-invalid BGP announcements to prevent the propagation of prefix origin hijacking.

Compared to route filtering based on IRR data, the data used by ROV can be validated cryptographically to ensure the accuracy. Compared to BGPsec [10], ROV achieves much lower computational overhead and reduces the impact on routing convergence. Therefore, ROV is considered the most promising solution to mitigate prefix origin hijacking.

C. Mutually Agreed Norms for Routing Security

The Mutually Agreed Norms for Routing Security (MANRS) is launched by ISOC in 2014. To improve the security of the Internet's global routing system, it encourages networks operators to deploy well-established routing security mechanisms. Specifically, it proposes four actions for network operators, with three Mandatory and one recommended:

- Action 1 (Mandatory): Prevent the propagation of illegitimate BGP announcements from customers.
- Action 2 (Recommended): Prevent traffic with spoofed source IP address.
- Action 3 (Mandatory): Enter contact information in IRRs or PeeringDB.
- Action 4 (Mandatory): Document intended routing announcements in IRRs or RPKI. Using IRRs is mandatory and using RPKI is recommended.

MANRS proposes Action 1 [18] because it believes filtering illegitimate routes from customers can constrain illegitimate routes to be inside the limited customer cone. It does not measure the effectiveness of deploying ROV in other directions or perform simulations. In this work, we try to understand ROV deployment in the real world and why MANRS Action 1 is not followed. Our survey in § V and simulation results in § VI find that practical deployments which do not follow MANRS Action 1 may even provide better protection against BGP hijacking.

D. Motivation & Methodology

Although MANRS has been advocating for network operators to at least deploy ROV at customer interfaces since 2014, recent studies reveal the harsh reality that many ASes are not following this action and still propagating RPKI-invalid prefixes.

To understand ROV deployment in real world and networks' compliance to MANRS Action 1, we first download BGP data from RouteViews [27] and RIPE RIS [28] to identify which ASes are propagating RPKI-invalid prefixes. Subsequently, we identify which ASes are accepting RPKI-invalid prefixes from their customers, providers, or peers, respectively. ASes that have accepted RPKI-invalid prefixes from at least one customer are considered not following MANRS Action 1. Then, we try to investigate whether network operators would be motivated to deploy ROV after receiving notifications. To this end, we present the first notification experiment to investigate the impact of notifications on the remediation rate of ROV. We further examine whether the use of different nudges (*e.g.*, baseline, social norms, authority, reminder, and elicitation) and native language in notification messages can provide more incentives for operators to remediate.

To understand why network operators are not following MANRS Action 1, we first conduct interviews with five network operators from different countries. We share the results of our measurement and notification experiment with them, and discuss with them the underlying reasons for non-compliant networks. After the interviews, we design a questionnaire and conduct a survey by sending emails to network operators and Network Operators Group (NOG) mailing lists. Our survey results reveal the business conflict between MANRS Action 1 and practical deployment, as well as other barriers that hinder ROV deployment in the real world.

To promote the further deployment of ROV, we first provide practical recommendations for deployment strategy. Consider deploying ROV at all external interfaces simultaneously is hard for many network operators, we conduct simulation experiments to evaluate the effectiveness of deploying ROV at different classes of interfaces. According to results of simulation experiments, we identify the most recommended ROV deployment strategy, which not only achieves the best effectiveness in mitigating prefix origin hijackings but also follows the business interests of network operators. We also propose recommendations for backup to reduce the risks of bugs in ROV implementation, and provide a vendor support list for networks to purchase equipment that supports ROV.

In the following, we introduce our measurement results in § III, elaborate our notification experiment in § IV, describe our survey results in § V, provide recommendations in § VI, outline future directions for research in § VII, compare with related work in § VIII, discuss limitations of our methodology in § IX, and finally conclude in § X. We summarize the abbreviations used in this paper in appendix § A.

III. Measurement

In this section, we first introduce the methodology of our measurement, and then describe the measurement results.

A. Measurement Methodology

RPKI-invalid prefixes propagation behavior. It is still a challenge to measure the deployment of ROV at scale, and none of prior researches can achieve a high level of confidence.

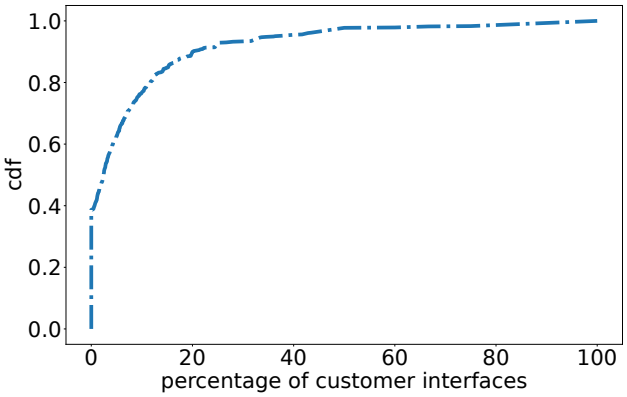

Figure 3: Percentage of customer interfaces that accept RPKI-invalid prefixes for non-stub ASes. More than 60% of non-stub ASes are not compliant to MANRS Action 1.

To ensure the measurement accuracy, we do not measure which ASes have deployed ROV, but focus on measuring the prevalence of RPKI-invalid prefixes that propagated through each AS. In other words, we only consider ASes that have passed RPKI-invalid prefixes. To this end, we try to detect as many ASes as possible by downloading data from all route collectors of RouteViews and RIPE RIS. The public route collectors peer with multiple vantage points (VPs), and collect BGP updates and Routing Information Bases (RIBs) from these VPs. We then perform ROV to determine the validation result of every BGP announcement. By checking the AS path of every RPKI-invalid BGP announcement, we can identify which ASes are propagating RPKI-invalid prefixes. This method may miss some ASes due to the limited number of vantage points, but it ensures the considered ASes do not deploy full RPKI-invalid filtering. We discuss this limitation in § IX. Finally, we identified 1,012 ASes that have propagated RPKI-invalid prefixes between June 15, 2022 and June 30, 2022, because the percentage of invalid prefixes is about 0.6% [29].

We further use the AS business relationship information provided by CAIDA [30], [31] to measure the prevalence of RPKI-invalid prefixes received from different classes of interfaces (*i.e.*, customer interfaces, provider interfaces, and peer interfaces). By check the filtering behavior at customer interfaces, we can identify which ASes are not following MANRS Action 1.

To quantify the prevalence of RPKI-invalid prefixes, we calculate the percentage of interfaces that accept RPKI-invalid prefixes for each AS using Formula (1):

$$P_{interfaces} = \frac{\#\ of\ interfaces\ that\ accept\ RPKI-invalid\ prefixes}{\#\ of\ interfaces} \quad (1)$$

and calculate the percentage of propagated RPKI-invalid prefixes for each AS using Formula (2):

$$P_{prefixes} = \frac{\#\ of\ propagated\ RPKI-invalid\ prefixes}{\#\ of\ propagated\ prefixes} \quad (2)$$

AS size. The routing complexity of an AS increases with the number of customers. ASes with more customers are more difficult to deploy ROV. To fairly compare the difference between ASes with different routing complexities, we classify ASes into three sizes (*i.e.*, small ASes, medium ASes, and large ASes) by following the thresholds defined in [32]:

- Small ASes: $number\ of\ customers\ \leq 2$
- Medium ASes: $2 < number\ of\ customers\ \leq\ 180$
- Large ASes: $number\ of\ customers > 180$

B. Measurement Results

As mentioned in § III-A, we finally find 1,012 ASes that have passed RPKI-invalid prefixes. Although the number is relatively small, the 1,012 ASes are distributed in different sizes. Therefore, our measure results help understand ROV deployment of ASes with different sizes. We also compare the difference between ASes with similar size.

Compliance to MANRS Action 1. Among the 1,012 ASes, 117 ASes are stub ASes, *i.e.*, the AS with no customer, and 895 ASes are non-stub ASes, *i.e.*, the AS with at least one customer. To understand AS-level compliance to MANRS Action 1, we calculate the percentage of customer interfaces that accept RPKI-invalid prefixes for each non-stub AS. Figure 3 shows the distribution of results. It shows that 61.3% of the 895 non-stub ASes have passed invalid prefixes received from customers. Since MANRS Action 1 requires ASes to deploy ROV at all customer interfaces, it means that the 61.3% ASes do not follow MANRS Action 1.

To further investigate whether ROV deployment is influenced by AS business relationships, we measure whether ASes have passed RPKI-invalid prefixes received from provider or peer interfaces. We find that 29.0% of the 1,012 ASes never accept any RPKI-invalid prefixes at provider interfaces, and 31.6% never accept any RPKI-invalid prefixes at peer interfaces. This suggests that some ASes may prefer to deploy ROV at provider or peer interfaces rather than at customer interfaces, which conflicts with MANRS Action 1.

In addition, our measurement indicates that 25% of the 1,012 ASes have passed RPKI-invalid prefixes from all classes of interfaces. The 25% ASes are more likely to not deploy any ROV.

Percentage of different classes of interfaces that accept RPKI-invalid prefixes. We try to investigate whether ASes with different routing complexities would use different deployment strategies. Figure 4(a) shows the distribution of the percentage of customer interfaces that accept RPKI-invalid prefixes for small ASes, medium ASes, and large ASes. It shows that 89.1% of small ASes accept no RPKI-invalid prefixes from customers, and only 39.0% of medium ASes and 2.3% of large ASes accept no RPKI-invalid prefixes from customers. Since about half of small ASes are stub ASes which have no customers, we further calculate the distribution for non-stub small ASes. As shown in Figure 4(a), there are up to 78.9% of non-stub small ASes propagating no RPKI-invalid prefixes received from customers. Figure 4(b) shows that 37.2% of large ASes and 31.7% of medium ASes accept no RPKI-invalid prefixes from providers, while only 18.0% of small ASes never accept RPKI-invalid prefixes from providers. Figure 4(c) shows that 36.1% of medium ASes accept no

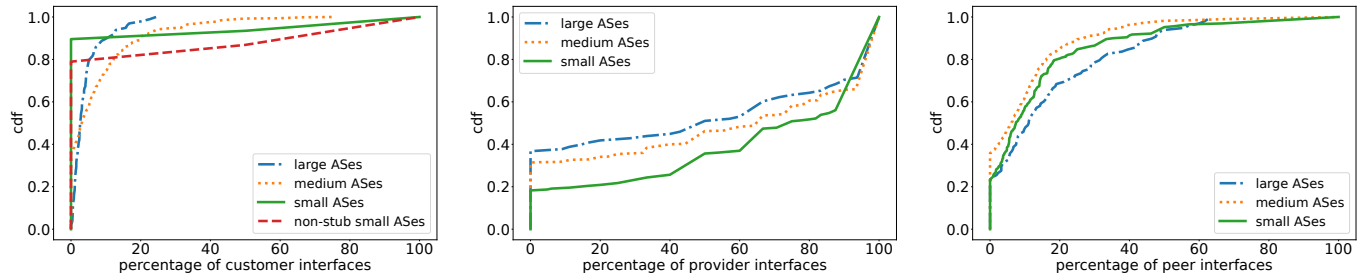

(a) Percentage of customer interfaces that accept RPKI-invalid prefixes.

(b) Percentage of provider interfaces that accept RPKI-invalid prefixes.

(c) Percentage of peer interfaces that accept RPKI-invalid prefixes.

Figure 4: Percentage of different classes of interfaces that accept RPKI-invalid prefixes.

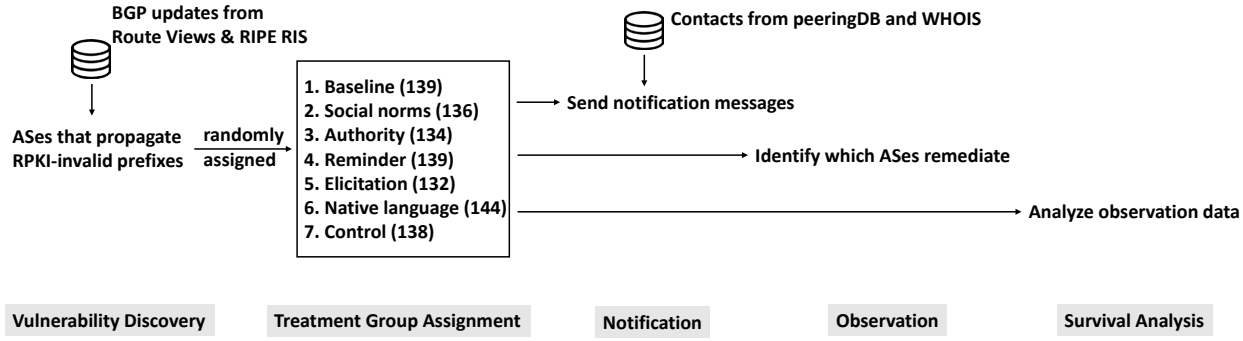

Figure 5: Overview of notification experiment.

RPKI-invalid prefixes from peers, compared to 20.9% for large ASes and 23.0% for small ASes.

We also compare the distribution of the percentage of propagated RPKI-invalid prefixes that received from different classes of interfaces between ASes with different routing complexities. The results are similar to the results shown in Figure 4.

Summary. Our measurement results show that more than 60% of ASes are not following MANRS Action 1, which requires networks to preferentially discard RPKI-invalid prefixes received from customers. Instead, some of them may choose to first deploy ROV at provider or peer interfaces. In addition, we find ASes with different routing complexities may prefer to perform different deployment strategies. More specifically, large ASes may be more likely to deploy ROV at provider interfaces than medium and small ASes. Medium ASes may be more likely to deploy ROV at peer interfaces than large and small ASes. While small ASes may be more likely to deploy ROV at customer interfaces than large and medium ASes. To understand the reasons for non-compliance and different deployment strategies, we conduct a survey in § V. Subsequently, we provide the most recommended deployment strategy according to our simulation experiments in § VI.

IV. Notification Experiment

In this section, we elaborate our notification experiment as well as our analysis results.

A. Overview

Figure 5 shows an overview of our notification experiment. There are five main procedures in this experiment. They are vulnerability discovery, treatment group assignment, notification, observation, and survival analysis.

First, we conduct a measurement to discover which ASes are propagating RPKI-invalid prefixes. Then, we randomly assign these ASes to 7 different experimental groups, including 6 treatment groups and 1 control group. To further test the effectiveness of different notification treatments, we design the unique message construction for each treatment group. After that, we send emails to ASes in treatment groups to notify their vulnerability to BGP hijackings and suggest them to deploy ROV in their networks. After sending notifications, we continue to observe which ASes take the remediation action in the next 55 days. Finally, we make a survival analysis based on the observation data to investigate the impact of notifications on ROV remediation.

B. Vulnerability Discovery

This procedure has been described in § III-A. We finally find 1,012 ASes that propagated RPKI-invalid prefixes between 15 June, 2022 and 30 June, 2022.

C. Treatment Group Assignment

Following the principle of randomized controlled trial [33], to compare the impact of notifications, we need to assign

the 1,012 ASes into treatment and control groups, and only send notifications to ASes in treatment groups. To further test the effectiveness of nudges and native language in notification messages, we specifically design 6 different treatment groups and 1 control group. Subsequently, we randomly assign the 1,012 ASes to the 7 groups.

Nudge Treatments. Nudge is an effective intervention that can help alter people's behavior in a predictable way [34]. Many behavioral science studies indicate that using nudges in the framing of messages can make people more likely to make a particular choice [35], [36], [37], [38]. Therefore, nudges are popularly employed in a wide variety of areas, including government policy, business management, healthcare, fundraising, and tourism [34]. We assume that nudges in the framing of notification messages can provide more incentive for recipients to remediate. To this end, we consider five nudge conditions in this notification experiment. They are normal notice of information (baseline), social norms, authority, reminder, and elicitation:

- *Baseline:* The first treatment group is called baseline group. In this group, we do not use nudges in the framing of notification messages, and notification messages are sent in English. In the text of baseline notification message[2], we use a wrong ASN or mismatched max-length example to explain that the recipient's network has propagated an RPKI-invalid prefix, and suggest the recipient to deploy ROV to protect its network from BGP prefix hijacking. In the following, we describe the other 4 nudge treatments, whose notification messages are designed based on the baseline notification message.
- *Social norms:* The social norms nudge is a widely used intervention strategy for promoting economic and public health behaviors. It raises collective behavioral expectations by emphasizing that other people in the social community have made or are preparing to make a particular choice. In the social norms nudge treatment, we point out that there are many networks that have deployed or are planning to deploy ROV, especially those participating in MANRS initiative. To this end, we add the following text to the baseline notification message: "Note that about 35% of network operators in the world have deployed ROV and the deployment ratio shows a positive trend. Particularly, most participants of MANRS promise to have deployed route filtering (with ROV as the best current practice)."
- *Authority:* Previous studies have found that people have a higher compliance with the recommendation suggested by an authoritative organization or institution than the recommendation suggested by an ordinary person. In the authority nudge treatment, we leverage the MANRS initiative to improve the authority of our notification. When sending notifications to networks in authority group, we add the following text to the baseline notification message:"Mutually Agreed Norms for Routing Security (MANRS), a global initiative supported by the Internet Society, has been calling on network operators to deploy route filtering mechanisms (with ROV as the best current practice) to secure the Internet."
- *Reminder:* People may put off making behavioral changes simply due to their inertia, procrastination, or forgetfulness. In this case, a reminder could play a very important role. In the reminder nudge treatment, after sending the initial notifications, we will wait one month and send a second round of notifications to networks that have not taken remedial action. In the reminder message, we add the following text to the baseline notification message: "In June 2022 and July 2022, We conducted two worldwide measurements of Route Origin Validation (ROV) deployment. We send this email to inform you that your network was found to propagate RPKI-invalid prefixes in the two measurements."
- *Elicitation:* Some studies report that people are more likely to participate in an activity if someone elicits their intention to carry it out. For example, a simple question about the future conduct can lead to a significant impact [36]. For this purpose, we design a short questionnaire to help recipients better understand ROV and elicit their intention to deploy ROV in the future. In the elicitation nudge treatment, we add the questionnaire link to the baseline notification message: "To better understand your concerns about BGP security, please complete the anonymous questionnaire:[LINK]."

Language Treatment. We also assume that using the native language in notification messages could attract recipients' attention, leading to a better compliance with our recommendation. Therefore, we set up a native language treatment group in which we send baseline notification messages in every recipient's native language. To eliminate the potential bias induced by native English speakers, we decide to focus only on differences between non-native English speakers in the two treatment groups.

D. Notification

Previous studies frequently encounter a high email bounce rate of over 50% [39], [40], [41], due to the incorrect or out-of-date contact information in WHOIS [42]. Recently, Lone *et al.* [43] propose that the reachability of notification emails can be effectively improved by prioritizing PeeringDB [44] and technical contacts. Therefore, we adopt the same method as Lone *et al.* in our notification experiment. To determine the appropriate contact for every non-deploying AS, we first check whether there is a technical email address in PeeringDB or WHOIS. If they correspond to two different technical email addresses, we prioritize the technical contact in PeeringDB because contacts in PeeringDB are considered more reliable [43], [45], [46]. If we cannot find a technical contact in PeeringDB and WHOIS, we choose to use the abuse contact and also prioritize the abuse contact in PeeringDB. Eventually, we determine the email address for every AS in the treatment groups.

On July 9, 2022, we removed 15 ASes (12 of which are in treatment groups and 3 are in control group) that have already remediated before our notification, and then sent notification messages to the remaining ASes in the 6 treatment groups. Of the 859 emails sent out, 49 emails (*i.e.*, 5.70%) are not successfully delivered. For the 49 undelivered emails, we try to send each notification message again with an alternative email address in PeeringDB or WHOIS. After that, 14 of the 49 emails are delivered. In the end, we successfully deliver a total of 824 notifications and receive an overall email bounce rate of 4.07%. The final number of ASes for each experimental

[2]The text of baseline notification message is presented in the appendix.

Table I: Relative risk ratios for different nudge treatments compared to the control group.

Group	Remediated	Exposed	RR	CI
Control	11	138	-	-
Baseline	15	139	1.35	[0.64, 2.84]
Social Norms	5	136	0.46	[0.16, 1.29]
Authority	13	134	1.22	[0.57, 2.62]
Reminder	13	139	1.17	[0.54, 2.53]
Elicitation	14	132	1.33	[0.63, 2.82]

group is shown in Figure 5. In the next 55 days (from July 9, 2022 to Sep. 1, 2022), we continue to observe which ASes turn to deploy ROV. For ASes assigned to the reminder treatment group, if they do not remediate after one month, we will send a second notification on Aug. 9, 2022.

E. Observation

After sending notifications on July 9, 2022, we continued to conduct weekly measurements for the deployment of ROV until Sep. 1, 2022. We use the method proposed by Galid *et al.* [21] to identify which ASes remediate. We weekly download BGP announcements from public route collectors and perform ROV on BGP announcements to get the validation results. Then, we seek an origin AS that originates both a non-invalid (*i.e.*, RPKI-valid or RPKI-unknown) route and an RPKI-invalid announcement. We check whether there is only one AS on the AS path of the non-invalid announcement that does not propagate any RPKI-invalid announcements, which means that this AS discards RPKI-invalid announcements but accepts non-invalid announcements. Following the process defined in [21], we determine that a non-deploying AS starts to take remediation actions if it does not propagate any RPKI-invalid prefixes and meets the above criterion for three different origin ASes. In this way, during the period from July 9, 2022 to Sep. 1, 2022, we have observed that 83 ASes in our notification experiment successively turn to deploy ROV.

Discussion: It is worth noting that, an AS meeting the above criterion is not necessarily due to the deployment of ROV. Other factors can affect the accuracy of the observation results. For example, a non-remediation AS may be mistaken for remediating if its upstream AS does not propagate RPKI-invalid routes to it due to the route policy or route filtering of the upstream AS, or if the RPKI-invalid routes that propagated through the AS are not be observed by the limited public route collectors. Nevertheless, we argue that the random assignment of our randomized controlled experiment can mitigate the influence of this problem. It is because that the accuracy problem affects treatment groups and control group almost equally, which means that we can still measure the impact of notifications by comparing the difference between treatment groups and the control group.

F. Survival Analysis

We then analyze the impact of different notification treatments on the remediation rate of ROV based on the remediation data observed from July 9, 2022 to Sep 1, 2022.

Nudge Treatments Analysis. To analyze the effectiveness of different nudge treatments, we first calculate the relative risk ratios (RRs) for each nudge treatment compared to the control group. The RR is the ratio of the probability of remediation in the treatment group versus the probability of remediation in the control group. In the control group, 11 of the initial 138 ASes have remediated until Sep 1, 2022. Taking the remediation ratio of control group as a benchmark, we compute the RR and the corresponding 95% confidence intervals (CI) for each nudge treatment group. If the CI includes 1, it means there is no significant differences in remediation rates between the treatment and the control. As shown in Table I, the CIs for all nudge treatment groups include 1, indicating that different nudge treatments groups have no significant impact on the remediation rate of ROV.

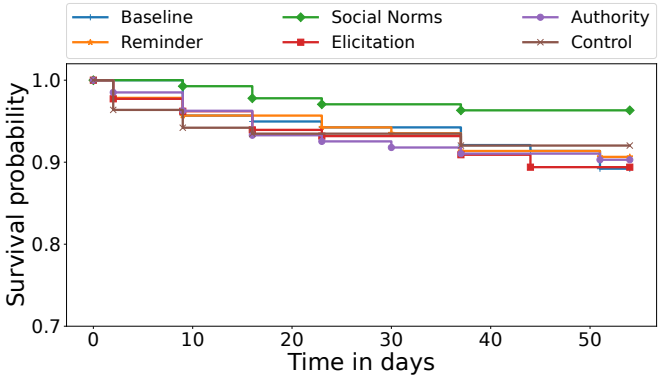

Figure 6: Survival curves for different nudge treatments and the control group.

Until the end of the observation period, some ASes in each group still do not deploy ROV and they may deploy ROV in the future. In this case, the results of RR analysis may have inaccuracy problems. Therefore, we further analyze the survival probabilities for different nudge treatment groups and the control group in statistic. We calculate Kaplan-Meier (KM) survival curves for each nudge treatment and the control. KM curve is the estimation of survival function $S(t)$, which is the probability that a subject survives longer than time t. Figure 6 shows the KM curves. The x-axis of Figure 6 is the time, from the beginning to the end of the observation period. The y-axis of Figure 6 is the probability of an AS not remediating t days after July 9, 2022. The downward trend of the KM survival curve is almost the same across all groups, except for the social norms nudge treatment. Intuitively, the downward trend of KM survival curve for social norms nudge treatment is slightly slower than other groups.

To check whether the survival function of each nudge treatment (especially the social norms nudge treatment) differs significantly with the control, we run the log-rank test for each nudge treatment compared to the control. Log-rank test is a hypothesis test to compare the efficacy of two treatments, which is the most commonly used test for data with censored observations. It tests the null hypothesis: $H_0 : S_1(t) = S_2(t)$ for all t, which means the two treatments are equally effective. If the result is $p \leq 0.05$, the differences between the two treatments are considered statistically significant. However, we find the significant value of every treatment compared to the control is greater than 0.05 (*i.e.*, $p > 0.05$). In particular,

the significance value of social norms nudge treatment is 0.12. Therefore, we can conclude that $none$ of the nudge treatments can significantly improve the remediation rate of ROV compared to the control group.

Native Language Treatment Analysis. We also analyze the impact of using native language in notification messages on remediation. Similarly, the survival analysis shows that there is no significant difference between the use of native language and English. More details can be found in appendix § C.

In our notification experiment, $none$ of the notification treatments can significantly improve the remediation rate of ROV, and most of ASes do not remediate in the 55-day observation period. In the following section, we try to understand why some networks are not willing to deploy ROV and why MANRS Action 1 is not followed.

V. Survey

In this section, we describe how we design the survey and what we learn from the survey results.

A. Survey Design

We first conducted interviews with 5 network operators from South Africa, Germany, China, and the Netherlands in March 2023. We introduced our measurement and survival analysis results to them, and discussed with them why MANRS Action 1 is not followed by many network operators.

Although the 5 network operators cannot be representative of the Internet, we learn a lot of valuable information from the interviews. We learn that business interest is one of the most important barriers for networks to perform RPKI-invalid filtering on BGP announcements received from customers. Some transit providers mention that their customer networks buy transit services from them and require them not to discard announcements of the customer networks. Therefore, they can not follow MANRS Action 1 to deploy ROV at customer interfaces. In addition, some reply that they are initially prepared to perform RPKI-invalid filtering at all classes of interfaces, but the lack of time delays their performing full filtering due to a large number of interfaces. While other interviewees that have already fully deployed ROV think RPKI-invalid filtering should be deployed at all classes of interfaces at the same time because they do not see any value in partial filtering.

Based on the results of our measurement, notification experiments, and interviews, we carefully design a questionnaire which consists of two main topics:

- Network operators' view on the deployment of ROV.
- Network operators' suggestions for improvements.

The full text of the questionnaire is presented in appendix § D. In brief, we ask network operators six main questions:

- Do you deploy or intend to deploy ROV at provider interfaces, customer interfaces, or peer interfaces?
- What are your reasons for not intending to deploy ROV at different classes of interfaces?
- Have you encountered any problems when operating ROV?
- Does the implementation guide provided by MANRS initiative provide effective assistance?
- What do you think are the priorities of deploying ROV at different classes of interfaces?
- What are your valuable experiences or suggestions for implementing or operating ROV?

To investigate the barriers of ROV deployment and solicit suggestions for improvements from the community, we launched an anonymous survey among ASes in our notification experiment from May 25 to June 25, 2023. We also sent the questionnaire to the mailing list of North American Network Operators Group (NANOG), Africa Network Operators Group (AFNOG), and South Asian Network Operators Group (SANOG). So far, we have received 82 responses in total.

B. Survey Results

Consider routing complexity may affect the cost and difficulty of ROV deployment, we calculate the distribution of the routing complexity for the 82 respondents. 45.1% of respondents have less than 10 customers, 26.8% have more than 10 but less than 100 customers, and the other 28.1% have more than 100 customers. For the number of providers, 54.9% of respondents have less than 2 providers, 26.8% have more than 2 but less than 10 providers, and 18.3% have more than 10 providers. For the number of peers, 13.4% of respondents have less than 10 peers, 50.0% have more than 10 but less than 100 peers, and 36.6% have more than 100 peers. According to the above statistics, we believe that our survey results should not be significantly biased.

The deployment strategy of RPKI-invalid filtering. Of the 82 respondents, 37.8% have performed RPKI-invalid filtering at all external interfaces, 17.1% have performed partial filtering, and the other 45.1% have not performed RPKI-invalid filtering yet. In other words, 62.2% of our respondents do not fully deploy ROV and may participate in propagating RPKI-invalid prefixes. We then ask these networks about their intentions to the remediation of ROV. 62.2% of them are planning to perform RPKI-invalid filtering in the future, but 24.4% do not plan to perform RPKI-invalid filtering. The other 13.4% have no clear intention.

Further, we ask them about their deployment strategy of RPKI-invalid filtering at different classes of interfaces. Their replies indicate that partial filtering can also occur even within the same class of interfaces. For RPKI-invalid filtering at customer interfaces, 63.4% of the respondents perform or intend to perform filtering at all customer interfaces and the other 36.6% are reluctant to follow MANRS Action 1. Specifically, 21.9% of the respondents only perform or intend to perform filtering at part of customer interfaces, and 14.6% do not intend to perform filtering at any customer interfaces. For RPKI-invalid filtering at provider interfaces, 45.1% of the respondents perform or intend to perform filtering at all provider interfaces, 18.3% perform or intend to perform filtering at part of provider interfaces, and up to 36.6% of our respondents do not intend to perform filtering at any provider interfaces. For RPKI-invalid filtering at peer interfaces, 54.9% of the respondents perform or intend to perform filtering at all peer interfaces, 18.3% perform or intend to perform filtering at part of peer interfaces, and 26.8% do not perform or intend to perform filtering at any peer interfaces.

Overall, we find that a larger proportion of respondents are intending to perform RPKI-invalid filtering at all customer interfaces than at all provider or peer interfaces. This may be due to the influence of MANRS initiative, which calls on network operators to at least deploy ROV at customer interfaces. In the following, we explain why many networks are still not compliant to MANRS Action 1 right now.

Reasons for non-compliance. We ask network operators about their reasons for not following MANRS Action 1. Although they all think BGP hijacking is a severe security threat, they are not compliant to MANRS Action 1 mainly due to the lack of time and effort, business conflicts, limited router capability, technical bugs in ROV implementation, and technical limitations in ROV mechanism.

43.9% of respondents claim that it takes time and effort to implement ROV, and they do not have sufficient time to implement ROV at so many interfaces. 23.2% do not perform RPKI-invalid filtering at all customer interfaces because some customers do not want their BGP announcements to be dropped. 15.6% report that they cannot deploy ROV with RPKI validator, because their equipment does not support the RPKI to Router (RTR) protocol [47] or the equipment's software does not meet the performance requirements. Unfortunately, they do not have sufficient money for equipment upgrade.

In addition to the economical reasons, 19.5% of respondents also concern that the deployment of ROV may affect the availability of inter-domain routing, since it is unclear if discarding RPKI-invalid BGP announcements would cause outage problems for their customers. 9.8% of respondents do not think ROV is effective to prevent BGP hijacking, especially when the hijacker also spoofs the origin ASN in the BGP announcement. 8.5% are not willing to deploy ROV at current stage, because they believe that ROV is fully effective only if most of the networks have deployed ROV. A few respondents also claim that they do not deploy ROV because it has been reported that current ROV is featured with a high false positive rate, resulting in many legitimate BGP announcements being discarded. One respondent points out another dilemma in deploying ROV: since it is unrealistic to operate validator by themselves, they need to leverage external validator, but it is hard to find a reliable external party.

Problems encountered while operating ROV. We also ask network operators about the problems they faced while operating ROV. 26.8% of the them report that if an intermediate/transit AS is polluted because it does not deploy ROV, even though they deploy ROV, their traffic can still be hijacked when the traffic passes through this polluted AS. 22.0% find that ROV cannot effectively identify most hijackings due to the limited adoption of ROA now. 19.5% complain that they have faced a lot of tricky bugs in the implementation of ROV. For example, bugs in RPKI RTR servers or router software can cause stale ROA data present in the router memory, resulting in legitimate BGP announcements being mistakenly discarded. Hence, a manual RTR session refresh is sometimes needed to remove stale ROA data.

In addition, in some routers' implementation of ROV, by default, the validation results of BGP announcements can affect BGP routing selection. Specifically, the routers preferentially select RPKI-valid BGP announcements over RPKI-unknown BGP announcements, and discard RPKI-invalid BGP announcements. Some network operators do not want their routers to perform this operation because it can lead to deoptimization in routing policies and induce a large number of unwanted BGP updates. However, 13.3% have trouble in bypassing this configuration in their equipment because the implementation is mandatory. Besides, 6.7% find that ROV significantly affects the capabilities and performance of their equipment. These problems can also hinder the deployment of ROV.

Suggestions for improvement. We further ask network operators whether the implementation guide provided by MANRS is sufficiently helpful for their deployment of ROV. Although more than 60% of respondents think that MANRS's implementation guide is effective in guiding network operators on how to deploy ROV, some also think that it is necessary to deploy ROV at all classes of interfaces instead of only at customer interfaces.

We then ask for their suggestions for improvements of RPKI and ROV, and they mainly ask for operation guide for ROV, correctness validation of the implementation for ROV in software, and route path authorization mechanisms. Some of them think that the documents provided by MANRS are not that useful for small networks, because it describes how to configure ROV, but lacks how to operate ROV (*e.g.*, troubleshooting), leaving inexperienced network operators helpless when faced with some difficulties. Some suggest that when operating ROV, you would better ensure that the validation results of BGP announcements do not affect the process of BGP routing selection. Otherwise, it can generate too many unnecessary BGP updates and even compromise other routing policies. To this end, they suggest that equipment vendors need to take this into account in their implementation for ROV, allowing users to enable or disable this function based on personal preferences. Several operators argue that ROV is not effective at preventing BGP path hijacking. Hijackers can easily evade the check of ROV by tampering with the origin ASN in the BGP announcement. Therefore, they urge the Internet community to provide the guide of route path authorization as soon as possible.

In addition to the suggestions for improving ROV, some respondents also ask for better registration mechanisms for legacy IP resources. They hope that RPKI could allow resource holders to sign ROA records for legacy IP resources with less conditions. Actually, ARIN has already allow legacy resource holders to sign ROA records with their legacy IP resources if they have signed the Legacy Registration Services Agreement (LRSA). However, a large number of legacy resource holders refuse to sign the LRSA because they think this agreement would take away some of their rights [48]. Since most of the legacy IP resources are in the ARIN region, many legacy resources are still not be covered by ROA now.

C. Survey Summary

The survey results answer the questions of why notification cannot significantly improve the remediation rate of ROV and why MANRS Action 1 is not widely followed in current stage. In summary, non-compliant networks are mainly due to economical and technical reasons.

Economical reasons. The economical reasons include lack of time and effort, business conflicts, limited router capability, and high operational overhead. The problem of lack of time and effort is the most common and relatively easy to solve. We believe that most of network operators would remediate after addressing the urgent work at hand, because 62.2% of respondents are planning to deploy ROV in the future. To save time in learning how to deploy ROV, network operators are strongly suggested to learn the necessary knowledge from the implementation guide provided by MANRS. The problems of business conflicts, limited router capability, and high operational overhead are relatively difficult to solve. For transit providers, propagating invalid announcements from customers will not harm themselves, but discarding customers' announcements means discarding money. In terms of the problem of limited router capability, network operators have to update the version or capability of their routers to support ROV, which requires high economic cost. In addition, even if the equipment is capable of performing ROV, the additional operational overhead associated with ROV may compromise the performance of the equipment.

Technical reasons. The technical reasons can be classified into two categories: technical bugs in the implementation of ROV and technical limitations of ROV mechanism.

Technical bugs. Technical bugs in RIR servers or router software can greatly compromise the effectiveness of ROV, and cause legitimate BGP announcements to be discarded. Addressing these problems is beyond the ability of network operators, so a large number of network operators do not remediate after receiving our notifications. Therefore, the equipment vendors and organizations that develop and maintain the implementation of ROV are required to carefully validate the correctness of their implementation.

Technical limitations. Technical limitations of ROV mechanism are considered the most challenging. For example, the effectiveness of ROV is extremely affected by the deployment ratio of ROA. Since many IP spaces are not covered in ROA records, ROV may identify the validation result of an illegitimate BGP announcement as "unknown" and thus improperly accepts it. What's worse, it is also frequently observed that ROV may mistakenly discard legitimate BGP announcements due to incomplete or misconfigured ROA records [49], [48]. Therefore, to improve the accuracy and effectiveness of ROV, it is necessary to increase the deployment rate of ROA. However, it is also a challenge to promote the deployment of ROA, because the deployment rate of ROV and the deployment rate of ROA are mutually affected. With a low deployment rate of ROV, operators have little incentive to deploy ROA, because IP resources signed in ROA records are protected by only a small number of ASes. And vice versa, with a low deployment rate of ROA, operators have little incentive to deploy ROV, because ROV is only fully effective when all IP spaces are covered in ROA records. Another challenging limitation is that deploying ROV in your own network only provides partial security. On the one hand, ROV fails to identify the BGP hijacking if a hijacker also spoofs the origin ASN in the AS path. The deployed network is still vulnerable to BGP path hijackings. On the other hand, preventing BGP hijacking cannot rely only on local filtering, but also depends on upstream filtering. Otherwise, the traffic originated from your network may still be forwarded to the hijacker by an upstream AS that does not perform route filtering.

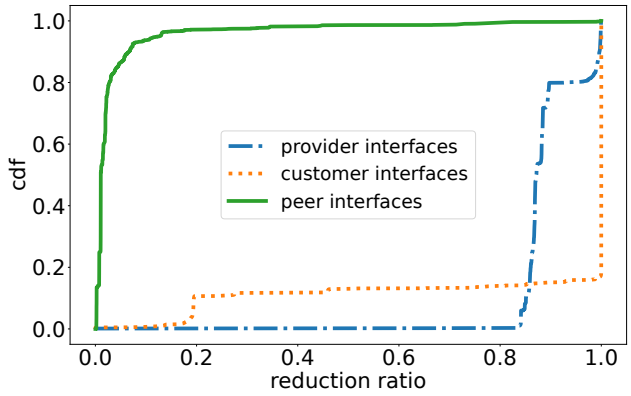

Figure 7: The distribution of reduction ratio of polluted ASes when all ASes deploy ROV at provider interfaces, at customer interfaces, or at peer interfaces.

VI. Recommendations

In this section, we provide practical recommendations that help improve the effectiveness of ROV and promote the deployment of ROV.

Summary of recommendations:

- We conduct simulation experiments to evaluate the priorities of deploying ROV at different classes of interfaces. We recommend that network operators could deploy ROV at provider interfaces as the first step to achieve the best effectiveness. This deployment strategy also solves the business conflicts between MANRS Action 1 and practical deployment.
- We provide recommendations for deployment backup, including geographic backup and software backup, to reduce the impact of technical bugs in ROV implementation.
- We identify a list of vendors that have achieved the implementation of ROV in their routers. Networks that need to upgrade equipment to deploy ROV are suggested to be purchased from these vendors.

A. Recommendation for Deployment Strategy

Consider it is difficult to perform RPKI-invalid filtering at all classes of interfaces simultaneously due to economical or technical reasons. Partial filtering is very common in the early days of ROV deployment for many networks. In the survey, we ask network operators about the priorities of deploying ROV at different classes of interfaces. However, there is no consensus among our respondents and some argue that following MANRS Action 1 may harm their business interests. To identify the best deployment strategy in the scenario of partial filtering, we conduct Internet-scale simulations and compare the effectiveness of deploying ROV at different classes of interfaces.

We first use the AS business relationship information provided by CAIDA [30], [31] to build the Internet topology with more than 74,000 ASes in our simulation experiments.

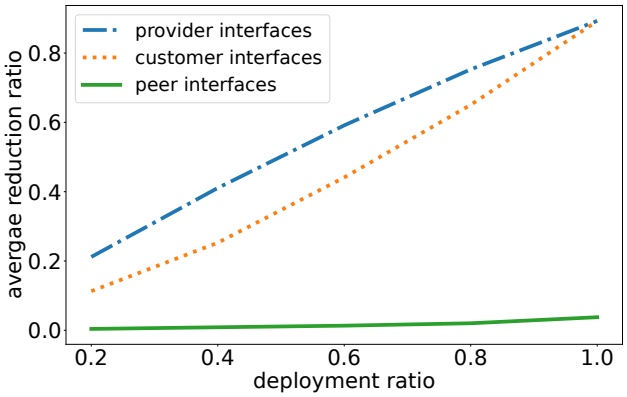

Figure 8: The average reduction ratio of polluted ASes of deploying ROV at provider interfaces, at customer interfaces, or at peer interfaces over different deployment ratios.

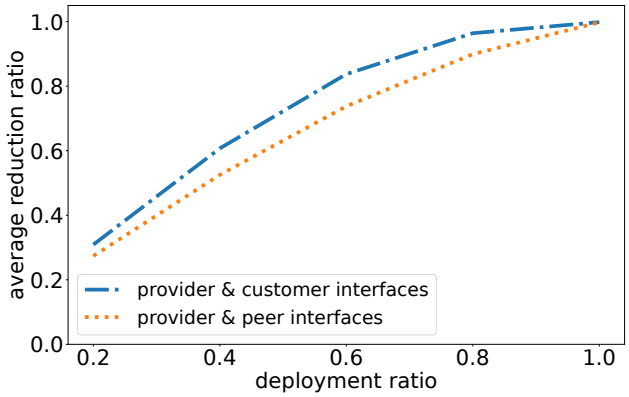

Figure 9: The average reduction ratio of polluted ASes of deploying ROV at provider and customer interfaces, or at provider and peer interfaces over different deployment ratios.

Then, we conduct 10,000 simulations on this topology. In each simulation, we randomly choose one prefix hijacker who announces an RPKI-invalid prefix in BGP. We simulate the propagation process of the RPKI-invalid prefix by implementing the routing tree algorithm [50] which considers valley-free principle. After that, we count the number of ASes that are polluted (i.e., accepting and propagating the RPKI-invalid prefix) in each simulation. To measure the effectiveness of deploying ROV at different interfaces, we enable RPKI-invalid filtering at all ASes' provider interfaces, customer interfaces, and peer interfaces, respectively. Under each of the three deployment strategies, we re-simulate the propagation process of the RPKI-invalid prefix in each simulation, identify how many previously polluted ASes now no longer accept and propagate the RPKI-invalid prefix, and finally calculate the reduction ratio of polluted ASes.

Figure 7 shows the distribution of reduction ratios in the 10,000 simulation experiments when all ASes deploy ROV at provider interfaces, at customer interfaces, or at peer interfaces. The x-axis shows the reduction ratio of polluted ASes, and the y-axis shows the corresponding cdf result. We find that ROV at peer interfaces has extremely limited reduction in the number of polluted ASes, and ROV at provider interfaces works best in preventing the propagation of RPKI-invalid prefixes. Specifically, when deploying ROV at all ASes' peer interfaces, for more than 90% of RPKI-invalid prefixes, the number of polluted ASes is reduced by less than 10%. In contrast, only 0.12% of RPKI-invalid prefixes have the reduction ratio of less than 10% when deploying ROV at all ASes' provider interfaces, and 0.67% when deploying ROV at all ASes' customer interfaces. In addition, the reduction ratio is more than 80% for 99.74% of RPKI-invalid prefixes when all ASes deploy ROV at provider interfaces, and for 86.01% of RPKI-invalid prefixes when all ASes deploy ROV at customer interfaces.

To measure the effectiveness of deploying ROV at different interfaces over different deployment ratios, we vary the deployment ratio of ROV-enabled ASes in each simulation from 20% to 100%. Figure 8 shows the average reduction ratios of the 10,000 simulations over different deployment ratios. As shown in Figure 8, deploying ROV at provider interfaces still achieves the best effect on reducing the number of polluted ASes under different deployment ratios, deploying ROV at customer interfaces is slightly inferior, and deploying ROV at peer interface has almost no effect. Particularly, when the deployment ratio is 40%, which is close to the deployment ratio in the real world, deploying ROV at provider interfaces can reduce the range of propagation of RPKI-invalid prefixes by an average of 41.1%, compared to 25.3% for deploying ROV at customer interfaces and 0.87% for deploying ROV at peer interfaces. Therefore, in the scenario of partial filtering, network operators are recommended to first deploy ROV at provider interfaces to achieve the best global effectiveness.

To identify the class of interfaces with the second highest priority, we further calculate the average reduction ratio of polluted ASes when deploying ROV at provider and customer interfaces, and at provider and peer interfaces over different deployment ratios. Figure 9 indicates that ROV at provider and customer interfaces works better in preventing the propagation of RPKI-invalid prefixes than ROV at provider and peer interfaces over different deployment ratios. When the deployment ratio is 40%, the combination of ROV deployment at provider and customer interfaces can reduce the range of propagation of RPKI-invalid prefixes by an average of 60.7%, compared to 52.5% for the combination of ROV deployment at provider and peer interfaces.

The most recommended deployment strategy. Based on the results of our simulation experiments, we recommend that networks that are unable to deploy ROV at all interfaces simultaneously could deploy ROV at provider interfaces as the first step, and deploy ROV at customer interfaces as the second step. Deploying ROV at peer interfaces could be done at the end. Following our recommended deployment strategy, the propagation of illegitimate BGP announcements can be best prevented under different deployment ratio scenarios. Moreover, for transit networks, deploying ROV at provider interfaces will not conflict with the business requirements of their customers. Therefore, we believe that the recommendation for deploying ROV at provider interfaces would encounter less economical resistance than MANRS Action 1. In this way, more non-deploying networks would be motivated to start performing partial filtering.

B. Recommendation for Backup

Since there may be potential bugs in RTR servers and software, we recommend network operators to validate the RTR servers and cache relying party software before performing route filtering based on RPKI. Moreover, to mitigate impact caused by possible bugs and outages, network operators are suggested to increase the geographic diversity and software diversity of ROV deployment. Specifically, they could deploy to two different data centers in case one has an outage and deploy two different code-bases in case one has a problem.

C. Recommendation for Purchasing

Consider some network operators reply that their routers do not support the implementation of ROV. We conduct market research to identify which equipment vendors have already supported the function of ROV in their router products. Based on the responses we received from different equipment vendors, we determine that Arista, Arrcus, Cisco, Extreme Networks, Huawei, H3C, Juniper, MikroTik, and Nokia have supported the function of ROV in their routers. Government agencies, organizations or individual network operators can refer to this list of vendors when purchasing new equipment. To address the expensive economic cost of updating equipment to support ROV, we agree that government agencies are encouraged to set up special grant programs to reduce the economic pressure on network operators to deploy ROV. To motivate more equipment vendors to support ROV in their equipment, government agencies can preferentially purchase equipment from vendors that support ROV.

In addition, in some vendors' ROV implementation, ROV results mandatorily affect the BGP routing selection. In our survey, as described in § V-B, some operators complained that the mandatory association would compromise their routing policy. To help address this problem and promote the deployment of ROV, we recommend that vendors' ROV implementation should allow users to enable or disable this function according to their needs.

VII. Future Directions

Since technical problems are difficult to solve in the short term, we particularly summarize a range of research directions in the future.

A. Automated Configuration and Operations

Correctly enabling and operating RPKI continues to pain network operators, particularly the smaller ones that lack technical knowledge and operational experience. Therefore, we believe it is worthwhile to investigate automated systems to assist operators in crafting correct and appropriate configurations for their routers, as well as in operations. The key challenges are that: (1) different vendors have different configuration syntax and vocabulary for their routers, and (2) to produce correct configurations, an automation system must understand its surroundings, *e.g.* how its neighbouring routers are configured. A recent work in unification of network devices' operations manuals [51] has shed the first light on tackling the first challenge, but the second challenge remains unaddressed.

B. Correctness Validation for ROV Implementation

To make it easier for equipment vendors to identify potential bugs in their implementation for ROV, an authoritative system is required to verify whether the implementation strictly follows the standard procedures of ROV mechanism. By using the correctness validation system, equipment vendors can ensure that their equipment properly supports ROV and avoids possible technical defects.

C. Route Path Authorization

ROV cannot identify hijackings when the hijacker also uses a forged origin ASN in the BGP announcement. Therefore, it is an urgent need to propose a brand-new routing security mechanism or propose improvements of RPKI to achieve route path authorization. BGPsec [10] can achieve route path authorization, but it is not widely used due to its high computational overhead. More recently, a more lightweight mechanism, Autonomous System Provider Authorization (ASPA) [52], has been proposed in IETF sidrops working group. An ASPA is designed as another kind of signed object in RPKI, which authorizes upstream providers for the customer AS. So, it can provide protection against some routing path hijackings. However, it is under discussion and has not yet become a standard. Therefore, the research of route path authorization is still in the preliminary stage of exploration.

D. Registration for Legacy IP Resources

Current RPKI architecture is not compatible with legacy IP resources, most of which are not included in existing ROA records. In the ARIN region, legacy resource holders are required to sign an additional LRSA agreement before they can be authorized to sign ROA records with legacy IP resource. However, a number of operators tell us that they are reluctant to sign the LRSA because they believe it takes away some of their rights. Therefore, a more friendly and less constrained registration mechanism for legacy IP resources is worth investigating. Ideally, this mechanism is supposed to be well integrated with the architecture of RPKI.

VIII. Related Work

A. ROV Measurement Methods

Existing ROV measurement methods [21], [53], [22], [54], [55], [56], [23] use control-plane route information, data-plane reachability information, or both to infer the deployment of ROV.

Galid *et al.* [21] present the first control-plane measurement of ROV deployment. They collect BGP announcements from multiple RouteViews collectors and implement ROV to check the validation result for each BGP announcement. ASes that propagate invalid prefixes are considered not to deploy ROV. ASes that do not propagate invalid prefixes but propagate valid or unknown prefixes for some origin ASes are considered to deploy ROV. On the basis of Galid *et al.*, Reuter *et al.* [53] and Gray *et al.* [22] perform more controlled measurements. They actively announce valid and invalid prefixes from PEERING testbed [57] and observe which ASes accept or drop the invalid prefixes. Particularly, Gray *et al.* further propose an algorithmic

framework by implementing Bayesian computation for ASes to determine whether an AS has deployed ROV.

Cartwright Cox *et al.* [54] and the RPKI WebTest [55] propose the data-plane measurement methods. They first sign ROA records in RPKI and announce a valid prefix as well as an invalid prefix against the signed ROAs. They then try to trigger active hosts or request participants in the test to send packets to the valid prefix and invalid prefix, respectively. By comparing the reachability to different prefixes, they can identify whether the local network of each host/participant is deploying ROV. However, compared to control-plane methods, the measurement range of data-plane methods is relatively limited.

Hlavacek *et al.* [56] and ROV-MI [23] combine control-plane and data-plane information to increase the measurement range. They typically collect BGP announcements from public route collectors and perform traceroute probes for valid and invalid prefixes. By using additional data-plane probes, they can learn more inverted paths and significantly extend the range of measurement.

B. Notification Experiments

Previous notification experiments have investigated the effectiveness of notifications in different security areas, including web misconfigurations [58], [59], [41], amplification DDoS attack [60], [61], and source address spoofing [43], [62]. They notify operators of security vulnerabilities in their networks and advise them to deploy corresponding security mechanisms.

Zeng *et al.* [59] investigate whether sending security notifications can help motivate website owners to remediate HTTP misconfigurations. They also test the effectiveness of different languages, constructions, persuasive solutions, and subject lines in notification messages. By comparing the remediation rates between different treatment groups and the control group, they find that security notifications have a moderate impact on the remediation rate.

Cetin *et al.* [60] focus on the remediation for DDoS prevention. They try to explore a more effective notification mechanism by quarantining the vulnerable network until it fixes the problem that its Network Time Protocol (NTP) servers can be abused in amplification DDoS attacks. They observe that quarantined networks tend to achieve higher remediation rates, about 87%, even though networks can easily exit from the quarantine environment. By contrast, for networks that are not notified or quarantined, only about half of them remediate the vulnerability.

More recently, Lone *et al.* [43] perform a notification experiment to advise operators to deploy source address validation (SAV). They also measure the effectiveness of different nudges in notification. However, unlike earlier experiments, their analysis indicates that none of treatment groups perform better than the control group.

In this work, we present the first notification experiment to investigate the impact of notifications on the remediation rate of ROV. Disappointingly, we find that none of notification treatments can significantly promote the deployment of ROV.

C. ROV Surveys

Previous studies [21], [63], [3] also conduct surveys on the deployment of ROV. However, Galid *et al.* [21] only ask network operators whether they deploy ROV and do not delve into the reasons for not deploying ROV. Other studies try to understand the barriers of ROV deployment, but their questionnaires are relatively simple and can provide limited information. For example, none of them focus on topics about deployment strategies at different classes of interfaces or the conflict between MANRS Action 1 and practical deployment.

To obtain more valuable information than previous surveys, we carefully design the questionnaire based on the interviews with several network operators as well as the results of our measurement and survival analysis. Overall, to the best of our knowledge, this work is the first time to analyze the impact of notifications on ROV remediation and systematically investigate the main obstacles to non-compliance.

IX. Discussion

In this section, we discuss the limitations of this work, present future works, and discuss the ethics considerations.

A. Limitations

The measurement methodology used in § III may not find all ASes that have passed RPKI-invalid prefixes, because this methodology passively observes BGP data on the Internet but some non-deploying ASes or some RPKI-invalid prefixes may not be observed by the limited public route collectors. The advanced ROV measurement methods, such as ROV-MI [23], can identify more than 4,000 ASes that do not deploy ROV, because, as described in § VIII-A, they additionally use data-plane probes to learn more information. Although our measurement method finds 1,012 non-deploying ASes, we can ensure these ASes must have passed RPKI-invalid prefixes, thus guaranteeing the accuracy of our measurement and notification experiment. Even though, we agree that if more non-deploying ASes can be discovered, the results of measurement and notification experiment can be closer to reality. Since ROV-MI does not publish its measurement results and model details, we will try to reproduce ROV-MI in the future.

In addition, we use the ROV measurement method proposed by Galid *et al.* to identify which ASes remediate in the observation process of notification experiment. Although none of existing ROV measurement methods have a high level of confidence and sufficient validation, we argue that the accuracy issues of our measurement method should affect each experimental group equally in § IV-E. Therefore, the accuracy issues do not seriously affect the results of the survival analysis.

In the survey, we only receive replies from 82 respondents, which is a small sample size compared to the more than 70,000 ASes in the Internet. Even though, our respondents consist of networks with different route complexities, and the distribution of the route complexity of the respondents is relatively balanced. Therefore, we believe that the respondents can represent ASes of different ranks without significant bias.

B. Future Works

To identify more ASes that are propagating RPKI-invalid prefixes, we plan to download BGP announcements from BGP monitors provided by more companies and organizations. In addition, since RPKI-invalid announcements make up a small proportion of BGP announcements, we plan to proactively announce some RPKI-invalid prefixes that are only used for experiments, and then observe which ASes are polluted by these RPKI-invalid prefixes.

We also plan to use more kinds of ROV measurement methods (such as ROV-MI) in the process of measurement and observation. We can combine the results of different methods to determine which networks are most likely to have fully deployed ROV, which are most likely to have deployed ROV at part of interfaces, and which are most likely to have not deployed ROV at all.

Consider the small sample size for our survey, we plan to have more in-depth interviews with a number of active network operators in the future. We decide to summarize the economical and technical problems in as much detail as possible and to classify them as urgent or trivial. This would help direct limited effort and resources to those urgent problems. We also plan to discuss these problems on more international platforms and seek help from more professionals.

C. Ethics Considerations.

To address possible ethical issues, we have consulted with our academic committee and department to ensure that our survey follows ethical principles [64] and protects participants' anonymity and rights to withdraw their answers. We strictly guarantee the anonymity of participants in our experiments and survey, and do not disclose their sensitive information in this paper. We have also sent emails to confirm participants' willingness to make their answers public and allowed them to withdraw their answers. None of them asked to withdraw their answers after receiving our emails. Besides, the questionnaire system we use can guarantee data protection and can provide the function of data deletion [65], [66]. We have removed any sensitive data from the questionnaire system that may raise ethical issues after the analysis.

X. Conclusion

In this work, we understand ROV deployment in the real world by conducting measurement and notification experiment, and investigate why many networks are not following MANRS Action 1 by conducting interviews and a more large-scale survey. To improve MANRS Action 1 and promote ROV deployment, we conduct extensive simulations to determine the most recommended deployment strategy for ROV. Following our recommendations, BGP prefix origin hijacking can be best prevented and the problem of business conflicts can be better avoided.

Acknowledgements

We thank our shepherd and reviewers for their thoughtful comments. Dan Li is the corresponding author. This work is supported by the National Key Research and Development Program of China under Grant 2022YFB3104800.

References

[1] "Ripe ncc. 2008. youtube hijacking: A ripe ncc ris case study. https://www.ripe.net/publications/news/industry-developments/youtube-hijacking-a-ripe-ncc-ris-case-study."

[2] "What is a ddos attack? https://www.netscout.com/what-is-ddos/bgp-hijacking."

[3] "Routing security goes to washington, https://www.internetsociety.org/blog/2022/04/routing-security-goes-to-washington/."

[4] A. Ramachandran and N. Feamster, "Understanding the network-level behavior of spammers," in *Proceedings of the 2006 conference on Applications, technologies, architectures, and protocols for computer communications*, 2006, pp. 291–302.

[5] P.-A. Vervier, O. Thonnard, and M. Dacier, "Mind your blocks: On the stealthiness of malicious bgp hijacks." in *NDSS*, 2015.

[6] M. Apostolaki, A. Zohar, and L. Vanbever, "Hijacking bitcoin: Routing attacks on cryptocurrencies," in *2017 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2017, pp. 375–392.

[7] "Doug madory. 2018. bgp hijack of amazon dns to steal crypto currency. https://dyn.com/blog/bgp-hijack-of-amazon-dns-to-steal-crypto-currency/."

[8] L. Qin, D. Li, R. Li, and K. Wang, "Themis: Accelerating the detection of route origin hijacking by distinguishing legitimate and illegitimate MOAS," in *31st USENIX Security Symposium (USENIX Security 22)*. Boston, MA: USENIX Association, Aug. 2022, pp. 4509–4524. [Online]. Available: https://www.usenix.org/conference/usenixsecurity22/presentation/qin

[9] P. Sermpezis, V. Kotronis, P. Gigis, X. Dimitropoulos, D. Cicalese, A. King, and A. Dainotti, "Artemis: Neutralizing bgp hijacking within a minute," *IEEE/ACM Transactions on Networking*, vol. 26, no. 6, pp. 2471–2486, 2018.

[10] M. Lepinski and K. Sriram, "Bgpsec protocol specification," *Internet Engineering Task Force (IETF)*, 2017.

[11] J. Schlamp, R. Holz, Q. Jacquemart, G. Carle, and E. W. Biersack, "Heap: reliable assessment of bgp hijacking attacks," *IEEE Journal on Selected Areas in Communications*, vol. 34, no. 6, pp. 1849–1861, 2016.

[12] S. Yingying, L. Dan, C. Li, L. Qi, and L. Sitong, "drr: A decentralized, scalable, and auditable architecture for rpki repository." in *NDSS*.

[13] M. Lepinski and S. Kent, "An infrastructure to support secure internet routing," Tech. Rep., 2012.

[14] X. Shi, Y. Xiang, Z. Wang, X. Yin, and J. Wu, "Detecting prefix hijackings in the internet with argus," in *Proceedings of the 2012 Internet Measurement Conference*, 2012, pp. 15–28.

[15] "Bgp security in 2021. https://www.manrs.org/2022/02/bgp-security-in-2021/."

[16] "October 25: As212046 – mezon – hijacked 3786 prefixes. https://twitter.com/Qrator_Radar/status/1452587538489778180?cxt=HHwWiIC9zfGq0KgoAAAA."

[17] "Mutually agreed norms for routing security (manrs). https://www.manrs.org/."

[18] "Manrs: Network operator actions. https://www.manrs.org/netops/network-operator-actions/."

[19] "Internet routing registry (irr). http://www.irr.net/."

[20] R. Bush, "Origin validation operation based on the resource public key infrastructure (rpki)," *IETF RFC7115 (January 2014)*, 2014.

[21] Y. Gilad, A. Cohen, A. Herzberg, M. Schapira, and H. Shulman, "Are we there yet? on rpki's deployment and security," 2017.

[22] C. Gray, C. Mosig, R. Bush, C. Pelsser, M. Roughan, T. C. Schmidt, and M. Wahlisch, "Bgp beacons, network tomography, and bayesian computation to locate route flap damping," in *Proceedings of the ACM Internet Measurement Conference*, 2020, pp. 492–505.

[23] W. Chen, Z. Wang, D. Han, C. Duan, X. Yin, J. Yang, and X. Shi, "Rov-mi: Large-scale, accurate and efficient measurement of rov deployment."

[24] "85% of manrs members conformant to actions 1 and 4. https://www.manrs.org/2023/01/85-of-manrs-members-conformant/."

[25] "As relationships. https://www.caida.org/catalog/datasets/as-relationships/."

[26] "Help validate rov adoption measurements from rovista. https://labs.ripe.net/author/tijay-chung/help-validate-rov-adoption-measurements-from-rovista/."

[27] "The route views project. http://www.routeviews.org/routeviews/."

[28] "Ripe ris (routing information service). https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris/ris-raw-data."

[29] "Nist rpki monitor. https://rpki-monitor.antd.nist.gov/."

[30] "As rank. https://asrank.caida.org/."

[31] M. Luckie, B. Huffaker, A. Dhamdhere, V. Giotsas, and K. Claffy, "As relationships, customer cones, and validation," in *Proceedings of the 2013 conference on Internet measurement conference*, 2013, pp. 243–256.

[32] A. Dhamdhere and C. Dovrolis, "Twelve years in the evolution of the internet ecosystem," *IEEE/ACM Transactions on Networking*, vol. 19, no. 5, pp. 1420–1433, 2011.

[33] "Randomized controlled trial (rct). https://en.wikipedia.org/wiki/Randomized_controlled_trial."

[34] "Application of nudge theory. https://en.wikipedia.org/wiki/Nudge_theory#Application_of_theory."

[35] R. H. Thaler and C. R. Sunstein, *Nudge: Improving decisions about health, wealth, and happiness*. Penguin, 2009.

[36] C. R. Sunstein, "Nudging: a very short guide," *Journal of Consumer Policy*, vol. 37, no. 4, pp. 583–588, 2014.

[37] D. Hummel and A. Maedche, "How effective is nudging? a quantitative review on the effect sizes and limits of empirical nudging studies," *Journal of Behavioral and Experimental Economics*, vol. 80, pp. 47–58, 2019.

[38] R. H. Thaler and C. R. Sunstein, "Libertarian paternalism," *American economic review*, vol. 93, no. 2, pp. 175–179, 2003.

[39] O. Cetin, M. Hanif Jhaveri, C. Gañán, M. van Eeten, and T. Moore, "Understanding the role of sender reputation in abuse reporting and cleanup," *Journal of Cybersecurity*, vol. 2, no. 1, pp. 83–98, 2016.

[40] O. Cetin, C. Ganan, M. Korczynski, and M. van Eeten, "Make notifications great again: learning how to notify in the age of large-scale vulnerability scanning," in *Workshop on the Economics of Information Security (WEIS)*, 2017.

[41] B. Stock, G. Pellegrino, C. Rossow, M. Johns, and M. Backes, "Hey, you have a problem: On the feasibility of {Large-Scale} web vulnerability notification," in *25th USENIX Security Symposium (USENIX Security 16)*, 2016, pp. 1015–1032.

[42] "Whois. https://who.is/."

[43] Q. Lone, A. Frik, M. Luckie, M. Korczyński, M. van Eeten, and C. Ganán, "Deployment of source address validation by network operators: a randomized control trial," in *2022 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2022, pp. 2361–2378.

[44] "peerdingdb. https://www.peeringdb.com/."

[45] A. Lodhi, N. Larson, A. Dhamdhere, C. Dovrolis, and K. Claffy, "Using peeringdb to understand the peering ecosystem," *ACM SIGCOMM Computer Communication Review*, vol. 44, no. 2, pp. 20–27, 2014.

[46] T. Böttger, F. Cuadrado, and S. Uhlig, "Looking for hypergiants in peeringdb," *ACM SIGCOMM Computer Communication Review*, vol. 48, no. 3, pp. 13–19, 2018.

[47] R. Bush and R. Austein, "The resource public key infrastructure (rpki) to router protocol, version 1," Tech. Rep., 2017.

[48] "Resource public key infrastructure (rpki) technical analysis. https://www.icann.org/en/system/files/files/octo-014-02sep20-en.pdf."

[49] T. Chung, E. Aben, T. Bruijnzeels, B. Chandrasekaran, D. Choffnes, D. Levin, B. M. Maggs, A. Mislove, R. v. Rijswijk-Deij, J. Rula *et al.*, "Rpki is coming of age: a longitudinal study of rpki deployment and invalid route origins," in *Proceedings of the Internet Measurement Conference*, 2019, pp. 406–419.

[50] P. Gill, M. Schapira, and S. Goldberg, "Modeling on quicksand: Dealing with the scarcity of ground truth in interdomain routing data," *ACM SIGCOMM Computer Communication Review*, vol. 42, no. 1, pp. 40–46, 2012.

[51] H. Chen, Y. Miao, L. Chen, H. Sun, H. Xu, L. Liu, G. Zhang, and W. Wang, "Software-defined network assimilation: bridging the last mile towards centralized network configuration management with nassim," in *Proceedings of the ACM SIGCOMM 2022 Conference*, 2022, pp. 281–297.

[52] "draft-ietf-sidrops-aspa-verification-09. https://datatracker.ietf.org/doc/draft-ietf-sidrops-aspa-verification/."

[53] A. Reuter, R. Bush, I. Cunha, E. Katz-Bassett, T. C. Schmidt, and M. Wählisch, "Towards a rigorous methodology for measuring adoption of rpki route validation and filtering," *ACM SIGCOMM Computer Communication Review*, vol. 48, no. 1, pp. 19–27, 2018.

[54] "Are bgps security features working yet?. https://blog.benjojo.co.uk/post/are-bgps-security-features-working-yet-rpki."

[55] "Rpki test. https://labs.ripe.net/author/nathalie_nathalie/rpki-test/."

[56] T. Hlavacek, A. Herzberg, H. Shulman, and M. Waidner, "Practical experience: Methodologies for measuring route origin validation," in *2018 48th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. IEEE, 2018, pp. 634–641.

[57] B. Schlinker, K. Zarifis, I. Cunha, N. Feamster, and E. Katz-Bassett, "Peering: An as for us," in *Proceedings of the 13th ACM Workshop on Hot Topics in Networks*, 2014, pp. 1–7.

[58] F. Li, G. Ho, E. Kuan, Y. Niu, L. Ballard, K. Thomas, E. Bursztein, and V. Paxson, "Remedying web hijacking: Notification effectiveness and webmaster comprehension," in *Proceedings of the 25th International Conference on World Wide Web*, 2016, pp. 1009–1019.

[59] E. Zeng, F. Li, E. Stark, A. P. Felt, and P. Tabriz, "Fixing https misconfigurations at scale: An experiment with security notifications," 2019.

[60] O. Çetin, C. Gañán, L. Altena, S. Tajalizadehkhoob, and M. Van Eeten, "Tell me you fixed it: Evaluating vulnerability notifications via quarantine networks," in *2019 IEEE European Symposium on Security and Privacy (EuroS&P)*. IEEE, 2019, pp. 326–339.

[61] M. Kührer, T. Hupperich, C. Rossow, and T. Holz, "Exit from hell? reducing the impact of {Amplification}{DDoS} attacks," in *23rd USENIX Security Symposium (USENIX Security 14)*, 2014, pp. 111–125.

[62] 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.

[63] P. Sermpezis, V. Kotronis, A. Dainotti, and X. Dimitropoulos, "A survey among network operators on bgp prefix hijacking," *ACM SIGCOMM Computer Communication Review*, vol. 48, no. 1, pp. 64–69, 2018.

[64] M. Bailey, D. Dittrich, E. Kenneally, and D. Maughan, "The menlo report," *IEEE Security & Privacy*, vol. 10, no. 2, pp. 71–75, 2012.

[65] "Qualtrics: Data protection & privacy. https://www.qualtrics.com/support/survey-platform/getting-started/data-protection-privacy/."

[66] "Qualtrics: Data deletion. https://www.qualtrics.com/platform/security/data-deletion/."

APPENDIX

A. A Summary of Abbreviations

- AFNOG: Africa Network Operators Group
- AS: Autonomous System
- ASN: Autonomous Systems Number
- ASPA: Autonomous System Provider Authorization
- BGP: Border Gateway Protocol
- CAIDA: Center for Applied Internet Data Analysis
- CDF: Cumulative Distribution Function
- CI: Confidence Interval
- DDoS: Distributed Denial-of-Service
- IETF: Internet Engineering Task Force
- IRR: Internet Routing Registry
- ISOC: Internet Society
- KM: Kaplan-Meier
- LIR: Local Internet Registry
- LRSA: Legacy Registration Services Agreement

- MANRS: Mutually Agreed Norms for Routing Security
- NANOG: North American Network Operators Group
- NIR: National Internet Registry
- NOG: Network Operators Group
- NTP: Network Time Protocol
- PKI: Public Key Infrastructure
- RIB: Routing Information Base
- ROA: Route Origin Authorization
- ROV: Route Origin Validation
- RPKI: Resource Public Key Infrastructure
- RR: Risk Ratio
- RTR: RPKI to Router Protocol
- SANOF: South Asian Network Operators Group
- SAV: Source Address Validation
- VP: Vantage Point

B. Baseline Notification Message

TITLE:

AS X is vulnerable to BGP prefix hijacking

CONTENT:

We have conducted a worldwide measurement of Route Origin Validation (ROV) deployment.

Wrong ASN example: We have observed that your network may not deploy ROV. Here is a BGP announcement received from public BGP monitor: [BGP ANNOUNCEMENT]. By querying Route Origin Authorization (ROA), [PREFIX] is not owned by AS Y. But AS X received this wrong-ASN BGP announcement and propagated it. You can obtain and check the BGP announcement from: [DATA LINK]

Mismatched max-length example: We have observed that your network may not deploy ROV. Here is a BGP announcement received from public BGP monitor: [BGP ANNOUNCEMENT]. By querying Route Origin Authorization (ROA), the maximum prefix length that [PREFIX] can be announced by AS Y is [MAXLENGTH]. But AS X received this too-specific BGP announcement and propagated it. You can obtain and check the BGP announcement from: [DATA LINK]

We encourage you to deploy ROV to protect your network from BGP prefix hijacking, and this is the implementation guide: https://www.manrs.org/netops/guide/

If you have any questions, issues, or concerns, please send an email to us.

C. Native Language Treatment Analysis

We focus only on differences between non-native English speakers in baseline group and native language group. We determine the corresponding country for every AS in this group by using country information collected from RIR databases. We then translate emails into Portuguese, Spanish, German, French, Italian, Chinese and more than 20 other languages.

Table II shows the relative risk ratio for using native language compared to using English in notifications sent to non-native English speakers. Since the 95% CI includes 1, we conclude that there is no significant differences in the remediation rate of ROV between using English and using native language. We further calculate the KM survival curves for the two language treatments and run log-rank test. Figure 10 shows the KM survival curves. The result of log-rank test for the two treatments is 0.48 (*i.e.*, $p > 0.05$), which means that there is no significant difference between the use of native language and English in notification messages.

Table II: Relative risk ratio for using native language compared to using English in notifications sent to non-native English speakers.

Group	Remediated	Exposed	RR	CI
Baseline	9	72	-	-
Native language	7	78	0.72	[0.28, 1.83]

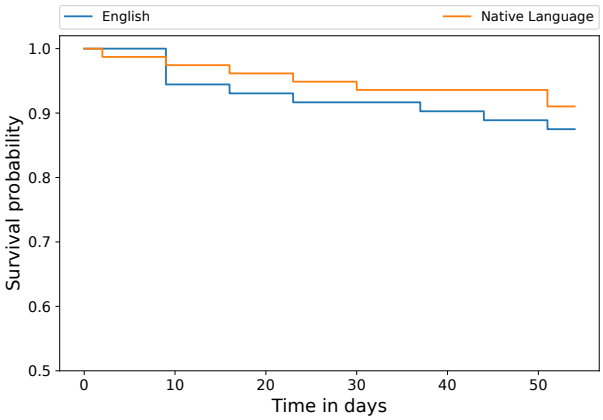

Figure 10: Survival curves for using native language and using English in notifications sent to non-native English speakers.

D. Survey Questionnaire

Q1: Have you performed Route Origin Validation (ROV) in your networks?

- Yes, validating BGP announcements received from all neighboring ASes
- Yes, but only validating BGP announcements received from part of neighboring ASes
- No
- Not sure

Q2: Are you planning to perform ROV in your networks?

- Yes
- No
- Not sure

Q3: How many customer ASes does your network have?

- 0-10
- 10-100
- More than 100
- Not sure

Q4: Do you perform or intend to perform RPKI-invalid filtering at all customer interfaces?

- Yes

- No, but I perform or intend to perform ROV only at partial customer interfaces
- No, I do not intend to perform ROV at any customer interfaces

Q5: How many provider ASes does your network have?

- 0-2
- 2-10
- More than 10
- Not sure

Q6: Do you perform or intend to perform RPKI-invalid filtering at all provider interfaces?

- Yes
- No, but I perform or intend to perform ROV only at partial provider interfaces
- No, I do not intend to perform ROV at any provider interfaces

Q7: How many peer ASes does your network have?

- 0-10
- 10-100
- More than 100
- Not sure

Q8: Do you perform or intend to perform RPKI-invalid filtering at all peer interfaces?

- Yes
- No, but I perform or intend to perform ROV only at partial peer interfaces
- No, I do not intend to perform ROV at any peer interfaces

Q9: What are your reasons for not performing ROV at all interfaces?

- Some adjacent ASes do not want their BGP announcements to be dropped
- I am concerned that implementing ROV may affect the performance of router
- I do not think ROV is an effective defense against BGP hijacking
- I do not think BGP prefix hijacking is a severe security threat
- It takes time and effort to implement and operate ROV at so many interfaces
- I think all BGP announcements received from customers should be accepted, even they are RPKI-invalid
- I think all BGP announcements received from providers should be accepted, even they are RPKI-invalid
- I think all BGP announcements received from peers should be accepted, even they are RPKI-invalid
- Others (please specify)

Q10: Have you encountered any problems since implementing ROV?

- It mistakenly drops legitimate BGP announcements
- It affects the capabilities and performance of routers
- It cannot effectively identify most hijackings due to the limited adoption of Route Origin Authorization (ROA) worldwide
- If an intermediate/transit AS is hijacked because it does not implement ROV, the traffic from my network will still be hijacked
- Others (please specify)

Q11: MANRS provides the implementation guide for ROV (https://www.manrs.org/netops/guide/). Do you think that is useful?

- Yes
- No
- Not sure

Q12: What do you think are the priorities of deploying ROV at different classes of interfaces? and please specify the reason.

Q13: Do you have any valuable experiences or suggestions for implementing or operating ROV? and please specify.