

Practical Experience: Methodologies for Measuring Route Origin Validation

*Tomas Hlavacek, Amir Herzberg, Haya Shulman, Michael Waidner
Fraunhofer Institute for Secure Information Technology
Darmstadt, Germany*

Abstract

Performing Route Origin Validation (ROV) to filter BGP announcements, which contradict Route Origin Authorizations (ROAs) is critical for protection against BGP prefix hijacks. Recent works quantified ROV enforcing Autonomous Systems (ASes) using control-plane experiments.

In this work we show that control-plane experiments do not provide accurate information about ROV-enforcing ASes. We devise data-plane approaches for evaluating ROV in the Internet and perform both control and data-plane experiments using different data acquisition sources. We analyze and correlate the results of our study to identify the number of ASes enforcing ROV, and hence protected with RPKI.

We perform simulations with the ROV-enforcing ASes that we identified, and find that their impact on the Internet security against prefix hijacks is negligible. As a countermeasure we provide recommendations how to cope with the main factor hindering wide adoption of ROV.

I. Introduction

The Internet's routing infrastructure with Border Gateway Protocol (BGP) has a long history of BGP prefix hijacks due to benign misconfigurations and malicious attacks, causing failures, outages and traffic interception, e.g. [1], [2], [3], [4], [5], [6].

BGP prefix hijacks have detrimental impact on stability and security of the Internet services and clients. The significance of BGP along with its insecurity generated multiple efforts to devise defenses. To allow networks to authenticate their prefixes the IETF designed and standardized Resource Public Key Infrastructure (RPKI) [7]. RPKI deployment consists of two complementary processes: Issuing Resource Origin Authorizations (ROAs) and performing Route Origin Validation (ROV). ROAs bind IP address blocks to owner ASes, listing which

ASes are authorized to originate routes for a given IP prefix. BGP routers then perform ROV using ROAs in order to filter routing announcements that violate the ROAs. This allows to detect and suppress IP prefix hijacks, in which an attacker or a misconfigured BGP router announces an IP address block that belongs to another AS. In addition to preventing prefix hijack attacks RPKI also forms a basis for other BGP security proposals, such as [8], [9], [10], for preventing more advanced and sophisticated attacks, such as *path manipulation*, [11].

Despite RPKI's importance for Internet security and extensive efforts to push its deployment forward, its adoption is progressing slowly and the networks are still exposed to traffic hijacks and outages due to misconfigurations and malicious attacks. Significant efforts are focused on understanding the security landscape of BGP and on evaluating the deployment of RPKI (i.e., the prefixes with ROA objects and ROV enforcing ASes), and identifying obstacles towards its wide adoption [12], [13], [14].

Quantifying prefixes for which ROA objects were created is easy, they are published in public RPKI repositories. Therefore, the number of ROAs and its growth rate are publicly known. Adoption of ROAs is progressing slowly yet there is a steady increase in the number of certified resources, most notably in RIPE NCC service region, which is leading the growth of public ROA repositories. Recent measurements show that about 6.5% of IP prefixes advertised in BGP are covered by ROAs [15], [16] but unfortunately, a large fraction of those ROAs are erroneous, [17], [18], [13]. Erroneous ROAs often cause ASes to lose legitimate traffic, hence demotivating enforcement of ROV filtering. How many ASes are performing ROV?

In contrast to measuring ROA growth rate, evaluating ROV adoption is a challenging task. How can one measure remotely whether some AS in the Internet performs ROV without having presence on that AS nor

having access to the routers on the network of the AS? A number of studies evaluated the implications of enforcing ROV locally on their own networks, [17], [18], [19], [14], and suggested that deploying ROV is likely to cause disruption of legitimate traffic. A recent study [13] examined ROV adoption in the Internet by a passive observation of existing BGP paths from multiple vantage points of 19 RouteViews [20] collectors and provided an upper bound on non-ROV enforcing ASes. The approach in [13] is based on observations of invalid prefix propagation in the Internet and measurements of ROV enforcement by leveraging valid BGP announcements as well as invalid ones, i.e., those that contradict ROAs. The study then checked for ASes that are on the paths towards the valid prefix but not on paths of the invalid prefix, hence concluding that those ASes filter invalid routes with ROV. Their findings showed that most large ASes did not deploy ROV and they estimated that 9 out of top 100 ASes could potentially be enforcing ROV. The study in [13] was based on *uncontrolled experiments* using control-plane data with only the BGP announcements that were generated by other ASes, none of which actively participated in the measurements. As a result the experiment could only provide partial information about the propagation path and even less information on the factors causing the propagation patterns, hence resulting in limited coverage and high false positives.

A subsequent study [21] argued that with *uncontrolled experiments* one cannot differentiate between the different causes for filtering of the invalid routes, and that often filtering of an invalid route is applied not due to ROV enforcement but also due to other factors, such as traffic engineering. To address this concern, [21] describes an approach for performing *controlled experiments* using BGP announcements and ROAs in combination with control-plane observations using RouteViews [20] and RIPE RIS [22] vantage points, and shows that the number of ASes performing ROV is smaller than the number that was found by [13].

Our work is inspired by these previous efforts to evaluate ROV adoption. We show that the previous methods do not accurately measure ROV deployment in the Internet, since these previous works solely relied on the BGP data from the (limited) set of RouteView collectors [20]. We perform an in depth evaluation by using, in addition to control-plane, also data-plane measurements. Furthermore, we use a significantly larger and more diverse set of collectors. Our results show significant differences, and indicate that complete reliance on the RouteView collectors can be quite misleading.

Unfortunately, our results indicate that ROV deployment has an even lower impact than is expected based on previous works. Our results show that a vast

majority of the ASes and of Internet users are not taking advantage of ROV and are hence vulnerable to prefix hijack attacks. This is demonstrated with our experiments where we propagate incorrect BGP announcements (contradicting ROAs) which are accepted by most of the networks.

Contributions. In this work we explore ROV adoption in the Internet. We perform controlled experiments and combine control and data plane routing information. Furthermore, in addition to the experiments performed in [13], [21] we perform two additional data collection experiments measuring real routing on the Internet, then compare and analyze the results. Our results show that there are discrepancies between the ROV percentage observed in vantage points and the real routing in the Internet, and that vantage points often do not reflect the accurate Internet routing but provide only a partial view. Our results portray a gloomy illustration of BGP security: *the number of ASes enforcing ROV is smaller than was found in [13] and even smaller than in [21]*. Using these results we perform simulations to demonstrate the scale of insecurity to prefix hijacks. We provide recommendations for coping with erroneous ROAs, to motivate more ASes to enforce ROV.

Organization. In Section II we present our methodology. We initiate with an active controlled control-plane experiment in Section III. Then in Sections IV and V we present our new approaches for ROV evaluation using data-plane, and describe the experimental evaluations based on them, along with the results. In Section VII we discuss the impact of the current ROV adoption on Internet security against BGP prefix hijacking and provide recommendations for countermeasures. We compare the results obtained from the three different approaches in Section VI and conclude this work in Section IX.

II. Evaluation Methodology

We devised two new approaches using data-plane for evaluating enforcement of ROV and performed two new experiments based on our approaches. The first approach (Section IV) is based on traceroute evaluations with RIPE Atlas and the second (Section V) is with TCP connection establishment to 1.25M-top Alexa web servers [23]. We also ran a control plane experiment (Section III) following an approach in [21] in order to validate our setup and results. Control-plane experiment with RouteViews and RIPE RIS is the only method that was used in prior work – it results in low noise and no bad paths. In RIPE Atlas *traceroute* and TCP probing experiments there are inverted paths (with difference in routing opposite to ROAs), furthermore, the measurements result in significant noise, hence require filtering out randomness from the data. Our work shows

that although control plane experiments are easier to launch, they are less accurate than our two new data-plane approaches and hence inferior. Specifically, the control-plane experiment produces results which are too positive and do not reflect the state of ROV deployment in the Internet.

Our study was performed over the period of five months, between February 2017 and June 2017, and is based on a controlled IPv4 prefix hijacking experiment. During all the experiments we announce two beacon prefixes $P_1 \equiv 188.227.158.0/24$ and $P_2 \equiv 188.227.159.0/24$. Both prefixes were announced from two distinct, geographically separated and meticulously selected Autonomous Systems (ASes) $A_1 \equiv AS29134$ and $A_2 \equiv AS378$. The prefixes P_1, P_2 and the Autonomous Systems Numbers (ASNs) A_1, A_2 are assigned to organizations that operate in RIPE NCC service region.

We created ROAs for the experimental prefixes and ASNs in RIPE hosted RPKI system in each evaluation batch according to the following table:

Measurement period	Published ROAs
February 2017 - May 2017	$\rho_1 = \{(P_1, A_1), (P_2, A_2)\}$
June 2017	$\rho_2 = \{(P_1, A_2), (P_2, A_1)\}$

We next describe the approaches and the results from the experimental evaluations.

III. Control-Plane Analysis

In this approach, used by [13] for an uncontrolled experiment, and later by [21] for a controlled experiment, valid and invalid (contradicting the ROAs) BGP announcements are propagated in the Internet. Analysis is then performed to check which routes the vantage points choose for the announced prefixes. We perform a similar experiment with active controlled manipulation of BGP announcements and of ROAs. This experiment, similarly to [13], [21], is limited to the control-plane observations. During the experiment we issue BGP announcements for prefixes we control; Figure 1 illustrates the elementary propagation pattern of our beacon prefix pairs in the Internet. The experiment uses two almost identical prefixes from the same origin, since both are likely to be propagated in the same way. Therefore, any differences in the propagation could indicate either manual interventions to the BGP best path selection algorithm or may be consequences of ROV filtering. To find the alteration in the path propagation we obtain and analyze BGP table dumps in MRT format ([RFC6396] [24]) from public BGP collectors - RIPE RIS and Route Views Project.

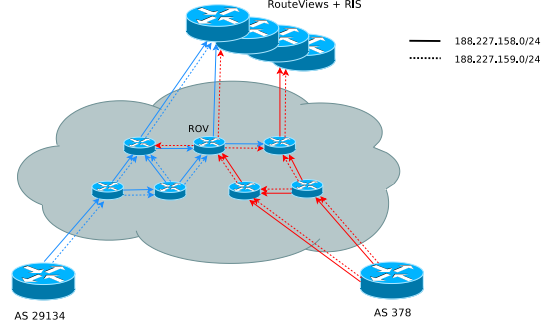


Fig. 1: Control-plane experiment

A. Experimental Evaluation

All available paths $X = \{\Pi_i | \forall i\}$ each consisting of prefix and AS-path: $\Pi_i = (p_i, \pi_i)$, are gathered from the vantage point BGP table dumps. The paths for our beacon prefixes P_1, P_2 : $X_{P_1, P_2} = \{\Pi_i | \forall i, p_i = P_1 \vee p_i = P_2\}$ are selected. AS-path $\pi = (a_1, a_2, \dots, a_n)$ from each path in X_{P_1, P_2} has to be obtained and categorized according to the route origin a_n compliance to the relevant ROA. The path can be either valid or invalid with respect to the currently published ROA. We define a validity symbol for path Π and ROA set ρ .

$$v(\Pi, \rho) = \begin{cases} 1 & \text{if } \Pi = (p, \pi), \pi = (a_1, a_2, \dots, a_n) \wedge \\ & \wedge (p, a_n) \in \rho \\ 0 & \text{otherwise} \end{cases}$$

All paths Π from X_{P_1, P_2} and the values $v(\Pi, \rho)$ were collected in both measurement time frames ρ_1 (February 2017 - May 2017) and ρ_2 (June 2017). This data yields groups of ASNs that falls into the following categories:

- (1) ASN a that occur in a path that is invalid with respect to a valid ROA: $\exists \Pi = (p, \pi), a \in \pi \wedge v(\Pi, \rho) = 0, \rho \in \{\rho_1, \rho_2\}$.
- (2) ASNs a that occur only in paths $\Pi = (p, \pi), a \in \pi$ that are compliant to ROA $v(\Pi, \rho)$ for $\rho \in \{\rho_1, \rho_2\}$.
- (3) ASNs a that satisfy the condition (2) and we have seen paths $\{(P_1, (\dots, a, \dots))^{\rho_1}, (P_2, (\dots, a, \dots))^{\rho_1}, (P_1, (\dots, a, \dots))^{\rho_2}, (P_2, (\dots, a, \dots))^{\rho_2}\} \subseteq X_{P_1, P_2}^{\rho_1, \rho_2}$ that traverses a for both prefixes P_1, P_2 in both time frames ρ_1, ρ_2 .
- (4) ASNs a that satisfy the condition (2) and we have seen at least two of four paths $(P_1, (\dots, a, \dots))^{\rho_1}, (P_2, (\dots, a, \dots))^{\rho_1}, (P_1, (\dots, a, \dots))^{\rho_2}, (P_2, (\dots, a, \dots))^{\rho_2}$ that traverses a for at least one prefix of P_1, P_2 in two different time frames ρ_1, ρ_2 .

The *first case* indicates that an AS a is not filtering out prefixes that fail ROA validation. Nonetheless, it is still possible that the AS could be de-prefering the non-compliant paths and the alternative compliant path was not available. The *second case* contains the ASes that do not transmit any non-compliant path in our experiment and could possibly filter or de-prefer

the non-compliant paths. This can be considered as an upper bound of the measurement. The *third case* is a subset of the second one: An AS qualifies for this category if we have evidence that it follows ROA for both prefixes in both time frames. Therefore the AS is very likely to filter BGP announcements according to ROV results or it can solely depend on an upstream AS that satisfies this condition. This category can be considered as a lower bound ROV acting AS. Manual analysis is then needed to remove false positives in ROV upstream dependents. The *fourth case* is a subset of the second and a superset of the third. It selects ASNs that we have no negative evidence about and there is a positive evidence that shows that the AS filters ROAs for at least one prefix in each measurement period. This category is defined because the requirements for the third one proved to be excessively strict. The third case requires the validating AS to receive the valid path for each prefix, which is not assured nor expected in all cases and therefore it generates false negatives. Thus we introduced the relaxed criteria (4) that helps us overcome this limitation.

The criteria are crafted to work with the two measurement rounds that are using the same prefixes and reversed ROAs to minimize possibility of false positives. There are still unpredictable factors, such as traffic engineering, that contribute to the path selection and certain combination can accidentally select valid paths. However, it is unlikely to happen twice in short time period and also to correctly react on reversing ROAs.

B. Data Analysis and Results

The presented results are consolidated and processed, and element counts of the previously explained groups are derived from RouteViews and RIPE RIS MRT dumps. The procedure involved downloading the MRT dumps after more than 24 hours after the first publication of ROAs and after injection of the prefixes from both origin sources and after the subsequent ROA changes. We repeated the analysis 3 times with variable delay among attempts and no significant difference has been detected in the entire result set.

The observed categories yielded following counts:

- Observed paths for both beacon prefixes: 696
- Total ASNs in the observed AS-paths: 296
- AS ROV categories:
 - (1) No validation (negative evidence): 250 (84.5%)
 - (2) Possible validation (upper bound): 46 (15.5%)
 - (3) Proved validation: 0
 - (4) Probable validation: 4 (1.35%)

The small scale and coverage limitation of this approach apparently affected the results and generated a low number of observed ASes. The reason lies in the

routing concentration in the Internet - most of the paths take route through a limited number of core ASes. Since the experiment was based on a fixed beginning of the paths, the only unique part in most of them was the ending of the path. However, the vantage points have in certain cases multiple peerings with one end ASN, which decreases the observation diversity of 696 paths to only 296 unique ASNs in combined AS-paths.

In contrast the coverage limitation could introduce positive bias. Arguably, the reason for obtaining favorable percentage of probable validating autonomous systems lies in the fact that the ASes that sponsor peerings with the vantage points count among the most progressive and technologically developed places in the Internet.

Moreover, control-plane analysis is affected by a known fundamental limitation: Route servers in most Internet Exchange Points (IXPs) do not include their ASN in AS-path and therefore a credit for dropping invalid prefixes might be given to the first AS downstream even though the filtering happens on the route server. In this case the ROV observation might be also lost if the downstream AS generates a negative evidence point as well.

IV. Traceroute Probes from RIPE Atlas

The most significant flaw of the control-plane analysis is limited visibility due to the low number of vantage points and thus insufficient Internet coverage. We strive to obtain representative data for the entire Internet. To employ more remote observation points we include a data-plane measurement. The method of choice for getting structurally similar data as in the control-plane measurement is remotely executed *traceroute*. RIPE Atlas [25] platform has been selected in order to run *traceroute* on several thousands remote points and collect the results in a uniform and machine-readable way.

A. Experimental Evaluation

Taking advantage of RIPE Atlas allows us to expand our experiment to more than 7.700 remote points. Figure 2 illustrates the setup of our experimental evaluation. Two consecutive *traceroute* runs are executed on each probe in our testing set. The first *traceroute* is directed to an IP address from the prefix P_1 and the second one to P_2 . Each *traceroute* result $\vec{T} \equiv (t_1, t_2, \dots, t_n)$ contains a lists of IP addresses that can be translated to a path $\Pi = (p, (A(t_1), A(t_2), \dots, A(t_n)))$ where p is the prefix of *traceroute* destination and $A(t)$ is symbol for resolving ASN for the router IP address t . With the derived path for each *traceroute* result we create a set of all obtained virtual paths X'_{P_1, P_2} and apply the reasoning from the previous section on these paths.

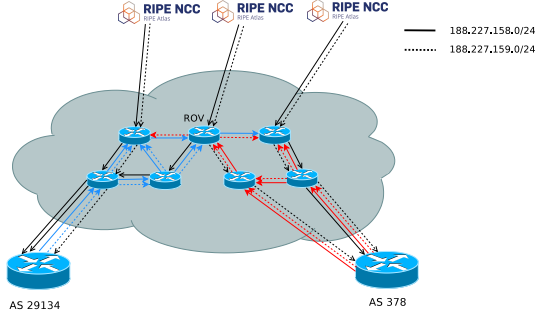


Fig. 2: RIPE Atlas *traceroute* experiment

The main difference lies in data acquisition and analysis method that has to be adapted for larger scale and for more significant noise level in traceroute data. However, the expected number of ASNs meeting the previously set criteria is still low. Thus, the results can be still manually checked, false positives removed and the positive results verified with the AS administrators or from external sources.

B. Data Analysis and Results

The RIPE Atlas results are convoluted from multiple *measurements*, that have to be scheduled separately due to RIPE Atlas restrictions on number of probes and maximum daily credit. Despite the limitations, it is still possible to emulate the control-plane measurement to a great extent. A new aspect in these measurements is the need for addressing data-plane noise, Atlas probe failures, misconfigured firewalls or other factors that cause incomplete or missing *traceroute* results. Figure 3 illustrates the noise effects, by plotting raw results of *traceroute* pairs comparison in a single measurement round.

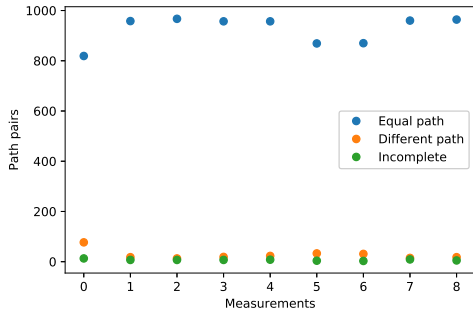


Fig. 3: Resolved Atlas pairs

The *traceroute* probe batches have been sent three times in each measurement period. The results were extracted, recombined to form the probe pairs and the *traceroute* results were transformed into derived virtual AS-paths for further categorization. The combined experiment statistics:

- Probes utilized in both time frames: 7730
- Identical path pairs: 7554 (97.7%)
- Different path pairs: 137 (1.8%)
- Incomplete pairs: 39 (0.5%)

The path pairs have been resolved to virtual AS-paths and the analysis identical to the control-plane measurement has been conducted. It yielded following AS ROV categories:

- (1) No validation (negative evidence): 2043 (97.0%)
- (2) Possible validation (upper bound): 49 (2.3%)
- (3) Proved validation: 2 (0.1%)
- (4) Probable validation: 12 (0.5%)

V. TCP SYN+ACK Capture

The control-plane measurement is limited by the low number of vantage points and low number of their peers, ranging in several hundreds. The RIPE Atlas measurement extended this range to a few thousand independent viewpoints. To extend the experiment further and measure ROV in larger portion of the Internet we perform active probing of servers.

A. Experimental Evaluation

The method is based on sending TCP connection initiation segments. Destination IP addresses are taken from 1.25M-top Alexa [23] top websites. The original list has been reduced to approximately 677K unique IP addresses suitable for the experiment. We call them d_1, d_2, \dots, d_N . Two distinct TCP SYN segment probes are sent to each destination d_i . A specific source IP address that lies in prefix P_1 is used for the first probe and an IP address from P_2 for the second probe. Our list of destinations contains mostly HTTP servers and therefore the probes are sent to a destination port 80/tcp to maximize the number of replies we get from the remote hosts. The TCP SYN+ACK replies from the destinations d_i are sent and routed to the probe source IP addresses in the prefixes P_1 and P_2 . We capture the TCP reply segments as they are routed towards the probe source IP addresses into the ASes A_1 or A_2 and we identify the probe destination, the reply receiving ASN and the probe source IP prefix. The desired coverage extension comes at a price of limited routing information that is extracted from each individual measurement. To reliably resolve a routing path that the reply packets take we have to adjust the objective of the experiment: We find destinations that benefit from the filtering provided by ROV to determine the number of protected sites in our destination set, instead of trying to identify individual ASes implementing ROV as in the two previous cases. Figure 4 shows the routing patterns for the probe replies. It follows the path propagation difference scheme and exploits the equal idea as in the previous two sections.

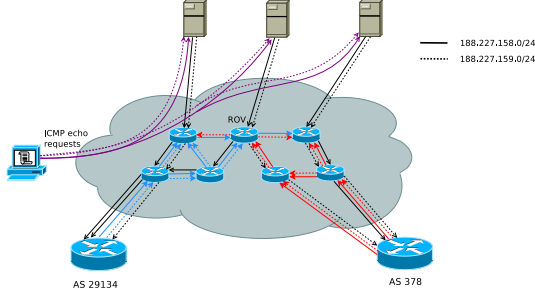


Fig. 4: ICMP reply capturing traceroute experiment

Even though we rely on the same attributes of BGP path selection algorithm as in the previous experiments, the structure and quality of data in this setup are completely different. The experiment output consist of pcap files that contain captured packets with the destination IP address lying in the prefixes P_1 and P_2 . The files were obtained from the ASes A_1 and A_2 in each measurement time frame ρ_1 and ρ_2 .

Let $R(d, P, A, \rho)$ be an indicator of response reception from d to an IP address in prefix P in AS A and assume that ROA set ρ was in effect at that time. We define symbols:

$$v'(d, A, \rho) = \begin{cases} 1 & \text{if } \forall P \in \{P_1, P_2\}, R(d, P, A, \rho) = 1 \Rightarrow (P, A) \in \rho \\ 0 & \text{otherwise} \end{cases}$$

$$\text{and } \bar{v}'(d, A, \rho) = \begin{cases} 1 & \text{if } \exists P \in \{P_1, P_2\}, R(d, P, A, \rho) = 1 \wedge (P, A) \notin \rho \\ 0 & \text{otherwise} \end{cases}$$

$$\text{and } v''(d, A, \rho) = \begin{cases} 1 & \text{if } \exists P \in \{P_1, P_2\}, R(d, P, A, \rho) = 1 \wedge (P, A) \in \rho \\ 0 & \text{otherwise} \end{cases}$$

The first symbol v' is an indicator of probe destinations d that in time frame ρ delivered into AS A only replies to an address in the correct prefix according to the valid ROA set. The second symbol \bar{v}' indicates destinations that at least once delivered a response to AS A , which contradicted ROA in time frame ρ . The last symbol v'' is an indicator of destinations d that delivered at least one reply in time frame ρ to AS A according to valid ROAs. Using these symbols we can define groups of destination IP addresses. The obvious definition for upper bound group is

$$\{d | \forall i, j \in \{1, 2\} : v'(d, A_i, \rho_j) = 1 \wedge \bar{v}'(d, A_i, \rho_j) = 0\}$$

However, this upper bound proves to be too permissive, because it contains all destinations that we have no negative evidence for. It includes even the destinations that do not respond to any of our probes. More precise requirements set yields the following sets:

- (1) $\{d | \exists A \in \{A_1, A_2\}, \exists \rho \in \{\rho_1, \rho_2\} : \bar{v}'(d, A, \rho) = 1\}$
- (2) $\{d | \forall \rho \in \{\rho_1, \rho_2\} \exists A \in \{A_1, A_2\} : v''(d, A, \rho) = 1 \wedge \bar{v}'(d, A, \rho) = 0\}$
- (3) Remaining destination IP addresses that do not fall into any of the previous groups.

The first group contains the prefixes that responded to the wrong origin AS at least once in our experiment. This group contains all destinations that are not protected by ROV. The second set contains the destinations for which we have at least two positive evidence points from two different time frames, hence this set contains the destinations that are likely to be protected by ROV. Destinations that do not respond at all, do not respond to most of our probes or responded randomly falls to the last group, that can be considered as unknown and unresolvable.

In this experiment a considerable noise has to be filtered or addressed otherwise. The distortion of observed routing can be caused by different factors ranging from load balancing techniques that produce random routing, traffic engineering, misconfiguration, network outages and ongoing changes in the Internet topology. To cancel out the effects of random packet drops we send multiple probes to our destinations. The complete measurement rounds have to be repeated several times with 24 hour time frame to spot random routing changes. And the results have to be combined together and fit into the simplistic scheme outlined by our group definitions.

B. Data Analysis and Results

The results from the TCP probing are more difficult to re-combine into pairs and post-process because of a high number of various errors generated by the data-plane. Figure 5 shows the number of lost replies and mismatched packets - replies that have been received from unknown sources, which were not queried, multiple replies to one probe and incorrect replies in general.

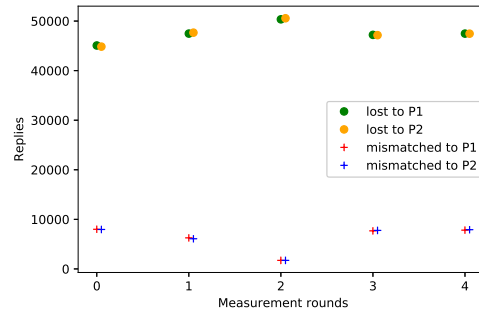


Fig. 5: Lost and mismatched TCP probes

Analysis of the pcap files according to the method described above yielded the following groups of the destination IP addresses:

- (1) Unprotected by ROV: 632570 (93.30%)

- (2) Likely protected: 201 (0.03%)
- (3) Unresolvable: 45163 (6.66%)

The results contain considerable number of unresolved destinations that failed to consistently respond to at least one of the ASes A_1 , A_2 . The probing batches consist of three uniform probes sent with approximately 3 second delays between them. Packet loss percentage for each destination has been determined from all the measurement batches and from both capturing points in A_1 and A_2 combined together. The destination was marked as unresolvable if it exhibited more than 90% packet loss for at least one prefix of the prefixes P_1 or P_2 , because absence of the negative evidence in these cases might be caused by packet loss rather than by ROA validation. Moreover, we observed considerable number of random routing cases (inconsistent routing behavior within or among probing batches), they are classified as *unprotected* in our experiment.

VI. Comparison of the Approaches

Figure 6 compares the results from three methods.

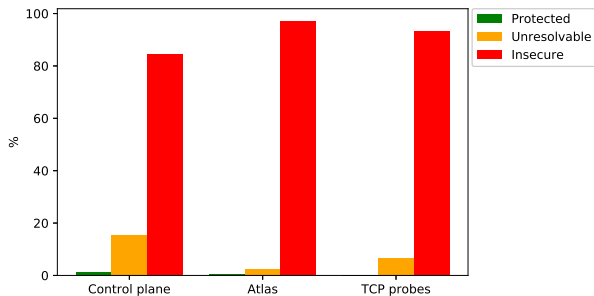


Fig. 6: Results from all measurement methods

Even though the numbers and resulting percentages can not be directly compared, following observations can be made: The higher percentage of protected entities found by the control-plane analysis and the Atlas *traceroute* experiment indicates that these measurement platforms are likely to introduce positive selection bias towards progressive technologies and better network management.

The active TCP probing experiment proved the expected fact, that Internet core autonomous systems do not deploy ROV. However, there are indisputable ROV deployments that have potential to protect certain portion of regional traffic. Unfortunately we were not able to find any large scale ROV deployment.

VII. Security Against BGP Prefix Hijacks

In this section we demonstrate the impact of current ROV adoption on the security of the Internet against

BGP prefix hijacks. We then provide recommendations for countermeasures.

A. Benefits from ROV Adoption

Our results indicate that very few ASes enforce ROV – less than the previous evaluations showed [13], [21]. What does this mean for Internet security against prefix hijacks? To answer this question we perform simulations on empirically derived datasets. Our simulations compute BGP routes using methods in [26], [27], [28] over the CAIDA AS-connectivity graph from December 2016. Our results average over 10^6 combinations of attacker and victim (i.e., the legitimate owner of IP prefix) ASes, both selected uniformly at random from the set of all ASes, as in [26], [28]. The results quantify to which extent the Internet is secure against BGP (sub)prefix hijacks given the ROV-supporting ASes that we found through our experimental evaluation in previous sections. Figure 7 shows projected upper bound of ROV adoption impact on number of autonomous systems that would not be affected by the simulated hijacking incidents and thus on routing security in the Internet.

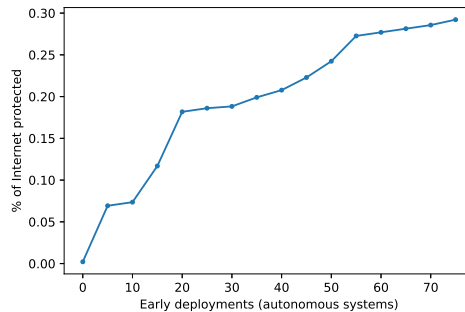


Fig. 7: ROV adoption impact on routing security

The percentage of protected AS is unfortunately low. Even the optimistic projection indicates that only a small portion of the prefixes in the Internet could benefit from ROV deployed in smaller ASes outside the Internet core. Our results illustrate an even more gloomy situation than [13] since they arbitrarily assigned probabilities for adoption of ROV, while we use our set containing ROV adopting ASes that we identified in our study. Generally, low ROV coverage in the Internet is not going to change unless the ASes at the core enforce ROV. Only wide spread adoption of ROV can help meeting the goal of considerable improvement in the Internet routing security.

B. Countermeasures

Given the state of security against prefix hijacks, adoption of ROV is paramount. However, as our results show, ROV adoption is essentially non-existent. What

can we do to improve adoption of ROV? One of the main obstacles towards wide enforcement of ROV is that network operators are concerned about loss of traffic if they apply ROV. This is mainly due to the many problematic ROAs, mostly due to errors in ROA issuance and to lack of coordination between different providers and customers. Our recommendation is a two-step solution: detect erroneous ROAs and then apply ROV only on ROAs that appear to be correct.

The first step can be performed using tools like `roalert.org`, which validates whether an AS has a valid ROA. In the second step the network can identify and then ignore the problematic ROAs and filter BGP announcements only according to ‘good’ ROAs.

VIII. Acknowledgements

We thank Hank Nussbacher, Israel Inter-University Computation Center for setting up the experiment and providing measured data.

IX. Conclusions

ROV enforcement by ASes is critical to protecting the Internet against BGP prefix hijacks. Previous efforts on measuring ROV adoption focused on uncontrolled [13] and controlled [21] control-plane experiments. We demonstrate that control-plane provides an estimate of the number and percentage of ROV adopters, however in reality the overall percentage is even lower. We provide two new approaches and describe our experiments based on them. We show that the percentage of adopters is smaller than found in previous research. We evaluated the ROV protection of the supporting ASes in face of BGP hijacking scenarios.

We demonstrated that the current ROV-enforcing ASes have only a negligible impact on the Internet security. We provide recommendations for tackling the main problem hindering wide deployment of ROV - the erroneous ROAs.

References

- [1] A. Toonk, “Hijack Event Today by Indosat,” <http://www.bgpmon.net/hijack-event-today-by-indosat/>.
- [2] “Renesys Blog - Pakistan Hijacks YouTube,” http://www.renysys.com/blog/2008/02/pakistan_hijacks_youtube_1.shtml, Feb. 2008.
- [3] A. Toonk, “Turkey Hijacking IP Addresses for Popular Global DNS Providers,” BGPMon.
- [4] “The New Threat: Targeted Internet Traffic Misdirection,” <http://www.renysys.com/2013/11/mitm-internet-hijacking/>.
- [5] H. Ballani, P. Francis, and X. Zhang, “A study of prefix hijacking and interception in the Internet,” in *SIGCOMM*, J. Murai and K. Cho, Eds. ACM, 2007, pp. 265–276. [Online]. Available: <http://doi.acm.org/10.1145/1282380>
- [6] P.-A. Vervier, O. Thonnard, and M. Dacier, “Mind Your Blocks: On the Stealthiness of Malicious BGP Hijacks,” in *NDSS*. The Internet Society, 2015. [Online]. Available: <http://www.internetsociety.org/events/ndss-symposium-2015>
- [7] M. Lepinski and S. Kent, “An Infrastructure to Support Secure Internet Routing,” RFC 6480 (Informational), RFC Editor, Fremont, CA, USA, pp. 1–24, Feb. 2012. [Online]. Available: <https://www.rfc-editor.org/rfc/rfc6480.txt>
- [8] R. White, “Deployment Considerations for Secure Origin BGP (soBGP).” June 2003. [Online]. Available: <http://tools.ietf.org/html/draft-white-sobgp-bgp-deployment-01>
- [9] A. Cohen, Y. Gilad, A. Herzberg, and M. Schapira, “Jumpstarting BGP Security with Path-End Validation,” in *SIGCOMM*. ACM, 2016, pp. 342–355. [Online]. Available: <http://doi.acm.org/10.1145/2934872>
- [10] M. Lepinski and K. Sriram, “BGPsec protocol specification,” September 2017, RFC8205. [Online]. Available: <http://tools.ietf.org/rfc/rfc8205.txt>
- [11] A. Cohen, Y. Gilad, A. Herzberg, and M. Schapira, “One hop for rpki, one giant leap for bgp security,” in *Proceedings of the 14th ACM Workshop on Hot Topics in Networks*, ser. HotNets-XIV. New York, NY, USA: ACM, 2015, pp. 10:1–10:7. [Online]. Available: <http://doi.acm.org/10.1145/2834050.2834078>
- [12] Y. Gilad, O. Sagga, and S. Goldberg, “Maxlength considered harmful to the RPKI,” in *CoNEXT*, 2017, pp. 101–107. [Online]. Available: <http://doi.acm.org/10.1145/3143361.3143363>
- [13] Y. Gilad, A. Cohen, A. Herzberg, M. Schapira, and H. Shulman, “Are we there yet? on rpki’s deployment and security.” NDSS, 2017.
- [14] D. Iamartino, C. Pelsser, and R. Bush, “boer,” in *PAM*, ser. Lecture Notes in Computer Science, J. Mirkovic and Y. Liu, Eds., vol. 8995. Springer, 2015, pp. 28–40. [Online]. Available: <http://dx.doi.org/10.1007/978-3-319-15509-8>
- [15] NIST, “RPKI Monitor,” <http://rpki-monitor.antd.nist.gov/>, 2015.
- [16] M. Wählisch, R. Schmidt, T. C. Schmidt, O. Maennel, S. Uhlig, and G. Tyson, “Ripki: The tragic story of rpki deployment in the web ecosystem,” in *Proceedings of the 14th ACM Workshop on Hot Topics in Networks, Philadelphia, PA, USA, November 16 - 17, 2015*, J. de Oliveira, J. Smith, K. J. Argyraki, and P. Levis, Eds. ACM, 2015, pp. 11:1–11:7. [Online]. Available: <http://dl.acm.org/citation.cfm?id=2834050>
- [17] J. Kloots, “RPKI Routing Policy Decision-Making: A SURFnet Perspective,” https://labs.ripe.net/Members/jac_kloots/, January 2014.
- [18] R. de Boer and J. de Koning, “BGP Origin Validation (RPKI),” Univeristy of Amsterdam, systems and network engineering group, Tech. Rep., July 2013.
- [19] D. Iamartino, “Study and Measurements of the RPKI Deployment,” 2015.
- [20] “University of Oregon Route Views Project,” <http://www.routeviews.org/>.
- [21] A. Reuter, R. Bush, Í. Cunha, E. Katz-Bassett, T. C. Schmidt, and M. Wählisch, “Towards a rigorous methodology for measuring adoption of RPKI route validation and filtering,” *CoRR*, vol. abs/1706.04263, 2017. [Online]. Available: <http://arxiv.org/abs/1706.04263>
- [22] “Routing Information Service (RIS),” <https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris>.
- [23] “Alexa Internet,” <https://www.alexa.com/>.
- [24] L. Blunk, M. Karir, and C. Labovitz, “Multi-Threaded Routing Toolkit (MRT) Routing Information Export Format,” RFC 6396 (Proposed Standard), RFC Editor, Fremont, CA, USA, pp. 1–33, Oct. 2011. [Online]. Available: <https://www.rfc-editor.org/rfc/rfc6396.txt>
- [25] “What is RIPE Atlas?” <https://atlas.ripe.net/about/>.
- [26] P. Gill, M. Schapira, and S. Goldberg, “Let the Market Drive Deployment: A Strategy for Transitioning to BGP Security,” in *SIGCOMM*, S. Keshav, J. Liebeherr, J. W. Byers, and J. C. Mogul, Eds. ACM, 2011, pp. 14–25.
- [27] —, “Modeling on Quicksand: Dealing with the Scarcity of Ground Truth in Interdomain Routing Data,” *Computer Communication Review*, vol. 42, no. 1, pp. 40–46, 2012.
- [28] S. Goldberg, M. Schapira, P. Hummon, and J. Rexford, “How Secure are Secure Interdomain Routing Protocols?” *Computer Networks*, vol. 70, pp. 260–287, 2014. [Online]. Available: <http://dx.doi.org/10.1016/j.comnet.2014.05.007>