

A tale of two synergies: Uncovering RPKI practices for RTBH at IXPs

Ioana Livadariu¹, Romain Fontugne², Amreesh Phokeer³, Massimo Candela⁴,
and Massimiliano Stucchi⁴

¹ Simula Metropolitan, Oslo, Norway
`ioana@simula.no`

² IIJ Research Laboratory, Tokyo, Japan
`romain@iiij.ad.jp`

³ Internet Society, Reston, USA
`phokeer@isoc.org`

⁴ NTT, Barneveld, Netherlands
`massimo@ntt.net`

⁵ AS58280, Brüttisellen, Switzerland
`max@stucchi.ch`

Abstract. Denial of Service (DoS) attacks and route hijacking have become the most predominant network attacks. To address these threats, network operators currently rely on mitigation services like Remotely Triggered Black Hole (RTBH) and Resource Public Key Infrastructure (RPKI). In this paper, we seek to understand how operators leverage both of these mechanisms. Using data collected at multiple IXPs we infer network operators that use RTBH services. We collect RPKI data for the same set of organizations and determine which of those rely on both RTBH and RPKI. One-third of the selected operators do not use any of these services, while most of the ASes that trigger blackholes also deploy RPKI. Some of these operators employ poor RPKI practices that make their prefixes vulnerable to attacks. However, most operators rely on an RTBH-agnostic approach indicating the need to devise an approach that effectively combines these two mechanisms.

1 Introduction

Distributed Denial of Service (DDoS) attacks are a pervasive threat to networking services and generally to the Internet Infrastructure itself. There have been different research studies to address the impact of DDoS attacks at scale [1–4], but in practice, they are difficult to tackle as they exploit the fault lines of the Internet, mainly due to the inherent lack of security at the protocol level.

Remotely Triggered Black Hole (RTBH) is a DDoS mitigation technique used to filter out undesirable traffic before it enters a protected network. BGP blackholing is signaled through BGP communities through “special” BGP announcements [5] that allow network operators to block unwanted traffic towards a specific destination. Network operators that rely on blackholing drop traffic

towards these targets, thus making them unreachable. Hence, this mitigation approach is highly effective against high-volume attacks. Many Internet Exchange Points (IXPs) have implemented RTBH to help protect participants' networks against unwanted traffic. However, as for any BGP announcement, blackhole routes are vulnerable to BGP hijacks such as prefix mis-origination and AS-path manipulation attacks. A solution to mitigate prefix mis-origination is to use Resource Public Key Infrastructure (RPKI) [6]. Using RPKI, network operators "authorize" specific Autonomous Systems (ASes) to originate their prefixes. This information is saved in the Route Origin Authorizations (ROAs) object which includes the prefix, the origin ASN, and the maximum length for the prefix.

In this paper, we analyze the prevalence of RPKI as a means to protect RTBH announcements against prefix mis-origination attacks but also we seek to understand the "challenging synergy" between these two frameworks. On one hand, RTBH is designed for fast⁶, sporadic, and very specific BGP announcements, while on the other hand, RPKI is designed to cover routes planned by operators, therefore avoiding overly specific prefixes (e.g., /32 for IPv4 and /128 for IPv6 address space). RTBH makes use of hyper-specific prefixes (greater than /24 in IPv4 and /48 in IPv6, or possibly even /32 in IPv4 and /128 in IPv6) and is very common at IXPs [8]. However, RFC 9319 recommends that operators SHOULD NOT create non-minimal ROAs (i.e., with max-length = /32 for IPv4 or /128 for IPv6) as it makes the announcement vulnerable to forged-origin sub-prefix hijacks and even discourages the use of RPKI to protect blackhole prefixes [9]. Currently, there are no mechanisms that combine the functionality of these two frameworks. Hence, network operators that use both systems have to find RPKI workarounds to ensure the RTBH effectiveness.

We seek to understand how operators use RPKI to protect RTBH announcements by collecting and analyzing BGP updates from 27 Internet Exchange Point (IXP) route servers [10]. We found that approximately 10% of the 2665 ASes that peer at the IXPs rely on the IXPs' blackholing service and most of the blackholed prefixes are /32s for IPv4 and /128s for IPv6. We also collected and analyzed RPKI data from the RIPE archive [11] to get a comprehensive view of the ROAs created by operators. Our results show that one-third of these operators do not deploy any of these two mechanisms, while most of the ASes that use RTBH also deploy RPKI.

Using these two datasets, we analyze how the RPKI management practices of the network operators impact the effectiveness of the RTBH announcements. Our analysis shows that few operators register very specific ROAs (/32 IPv4 and /128 IPv6) for IPs that are susceptible to DoS attacks. We also find that some operators create ROAs with a maximum value for the maxLength attribute. Networks that use this approach, however, can potentially have their IP address space hijacked [12]. However, most of the operators appear to rely on the IXPs' blackholing services which means that RTBH announcements from these members are going to be most likely rejected.

⁶ A recent study by Fontugne *et al.* [7] showed that RPKI can add significant delays to the propagation of BGP announcements.

2 Background and Related Work

2.1 RTBH and RPKI

Remotely Triggered Black Hole (RTBH). BGP blackholing is a common technique used to mitigate DoS attacks. Network operators that detect such attacks on specific targets in their network can rely on BGP updates to blackhole the traffic towards these targets and thus making these destinations unreachable. Hence, this mitigation approach is highly effective when the attack volume is high. IXPs are interconnection points in the Internet infrastructure that facilitate traffic exchange among a significant number of networks. Thus, IXPs are the perfect locations to implement blackholing services. In fact, many IXPs provide RTBH services to protect their members from unwanted traffic. During such an attack, the victim announces the target IP prefix using a well-known *BGP blackholing community*. Route servers propagate the RTBH announcement to other IXP members. After accepting the BGP update, the traffic destined for the victim is forwarded to the blackhole.

Resource Public Key Infrastructure (RPKI). The Resource Public Key Infrastructure (RPKI) [13, 6] is a public key infrastructure designed to protect operators against BGP hijacks. Network operators use RPKI to digitally sign Route Origin Authorization (ROA) objects, through which the IP address space owner states the *IP prefix(es)* and the *AS* number authorized to advertise the address block. ROAs also contain a *maxLength* attribute which allows the authorized AS to advertise multiple prefixes within the limits set by this attribute. Routers on the Internet use the information in RPKI (via a Relying Party) to validate incoming announcements through a process called Route Origin Validation (ROV). Based on the ROV outcome and on the router’s local policy, ingress BGP announcements are either allowed or dropped.

ROV by IXP route servers. RTBH was not designed with RPKI in mind and vice-versa, hence an unplanned blackholing announcement is likely to be RPKI invalid and dropped by networks implementing ROV. Therefore, in order to make RTBH and RPKI work together IXP route servers have to implement exceptions to handle these special cases. From the considered IXPs, DE-CIX is the only IXP that documents a detailed approach to implementing ROV and RTBH [14]. Equinix also lists how it implements RPKI [15]. Specifically, the IXPs route servers employ a loose ROV process that disregards the ROA’s *maxLength* attribute for matching BGP announcements marked with the blackholing community or next hop IP address. IXPs that rely on BIRD [16] can choose to implement the same approach as DE-CIX [17] for RTBH and ROV. Manually checking the IXP route servers in our study, we find that five route servers run BIRD [18–20]. Moreover, we devised and sent to IXPs a survey that aims to understand the current practices in RTBH and ROV. Note that we targeted the IXPs selected in this study. We received one reply that confirmed our assumption. This loose RPKI validation has the disadvantage of being vulnerable to *maxLength* attacks [12, 9]. Furthermore, although the route servers allow IXP members to make these, strictly speaking, RPKI invalid announcements to prop-

agate to other members, the members that implement ROV are likely to drop these announcements since the loose ROV process is deprecated for ASes.

2.2 Related Work

During the last few years blackholing has been the focus of several research studies [21, 3, 22–25]. Streibelt *et al.* showed how the propagation of BGP communities can be exploited by third-party networks to trigger remote blackholing [25]. Focusing on the Internet-wide usage of blackholing, Giotsas *et al.* found an increase in the BGP blackholing activity [21]. The authors inferred RTBH services from a large number of transit providers and a few IXPs. In their work [3] Dietzel *et al.* analyze the usage of the blackholing service of one IXP over the span of three months and report heavy usage of this service. The same authors later proposed an advanced blackholing system [22]. Wichtlhuber *et al.* also proposed a machine learning-based system for detecting and filtering DDoS attacks at IXP that considers as input blackholing announcements [24].

The Internet Engineering Task Force (IETF) has placed significant efforts in developing the RPKI management system [13, 6]. Researchers and network operators have started studying different aspects of the RPKI deployment [12, 26, 27, 7]. In our work, we seek to evaluate how RPKI and RTBH can efficiently be utilized by network operators. A recent IETF draft [28] looked into this question and proposed an approach to combine both RPKI and RTBH. Thus, indicating the need to understand how these two systems can efficiently co-exist. To the best of our knowledge, no other measurement study has focused this problem.

3 Blackholing activity at IXPs

3.1 BGP data from RTBH-enabled IXPs

We collect and analyze BGP data from Packet Clearing House (PCH) over the span of four months (from April to July 2022) to identify the prefixes blackholed by network operators. PCH maintains 217 route collectors hosted at different IXPs and located in 165 different cities across 95 countries [29]. In our study, we manually identify the PCH IXP route servers that publicly documented RTBH services. Thus, our collected BGP data comes from seven IXPs that are monitored by 27 PCH collectors located in Europe (16), United States (10) and Australia (1). Previous studies have extensively analyzed blackholing activity at IXPs [22, 3]. However, these studies focused only on one location, whereas our study includes a higher number of route servers. Moreover, our study does not solely focus on RTBH practices at the selected PCH route collectors but aims to understand the synergies between two network attack mitigation mechanisms.

The collected data includes BGP announcements originated by both the IXP members and their customers. In order to avoid wrong inferences due to BGP communities that are unintentionally propagating across multiple ASes [25] we filter out all BGP announcements originated by ASes that are not members of

the IXP focusing in this study only on the IXP members. Consequently, our BGP dataset contains 239,345 IPv4 and 79,162 IPv6 prefixes advertised by 2665 ASes. These numbers are not uniformly distributed across the PCH route collectors, DE-CIX Frankfurt, Equinix New York, and France-IX Paris have the highest number of ASes originating prefixes.

3.2 Identifying blackholed prefixes

To identify when a blackhole is remotely triggered, we leverage the RTBH documentation provided by IXPs [30–34]. Members of the selected IXPs can blackhole traffic by crafting BGP announcements that either contain a specific BGP community and/or a specific next-hop IP address.

For most of the selected IXPs, this community is 65535:666, while the next-hop IP address varies from one IXP to another. We monitor these attributes in our BGP dataset and set the *start of a blackholing period* for a prefix when we first observe the blackholing community or next-hop IP address prefix BGP announcement information. The *end of the blackholing period* corresponds either to the blackholed prefix withdrawal or the prefix being re-announced without any blackholing signal (i.e., BGP community or next-hop IP address).

We identify 12,670 IPv4 and 32 IPv6 prefixes blackholed by 225 members across 24 (out of the 27) analyzed collectors. We compute the number of such prefixes at each monitor and plot in figure 1 these values. Our analysis also shows variability in the blackholing activity across the collectors. Significant part of prefixes are /32s for IPv4 and /128s for IPv6. For half of the collectors analyzed, we observe more than 400 IPv4 blackholed prefixes. We also found that some locations have more blackholing activities, for example, more than 10% of all prefixes ob-

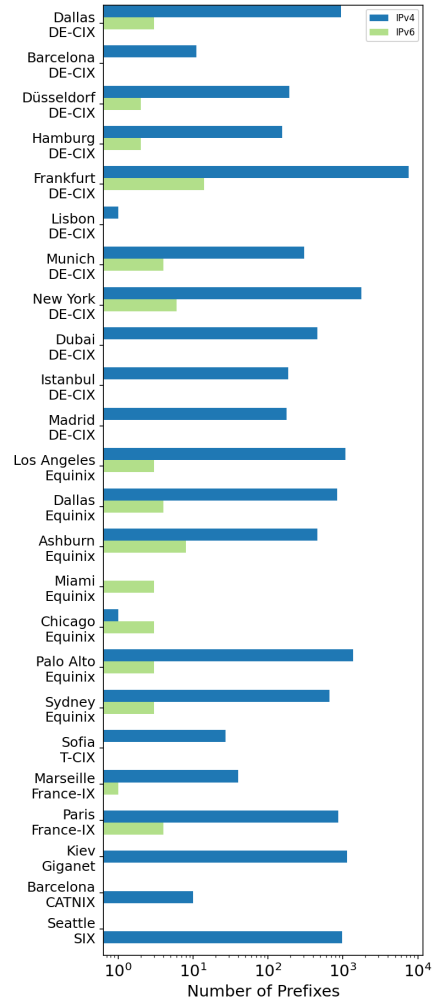


Fig. 1: Total number of unique blackholed prefixes at each of the 24 BGP collectors from April to July 2022.





Fig. 2: Distribution of the number of synchronized /32s per /24 prefixes per collector. The figure shows only collectors with more than 150 blackholed IPv4 /32 prefixes.

served at DE-CIX Frankfurt and Giganet Kiev are blackholed prefixes. We hypothesize that DE-CIX Frankfurt is a favored place for mitigating DDoS traffic because it is one of the largest IXPs, and the relatively high number of blackholed prefixes observed at Giganet Kiev is most likely due to the ongoing conflict in Ukraine [35–38].

Focusing on the ASes that trigger these blackholes, we observe an average of 57 IPv4 and 2 IPv6 prefixes per AS. Only one-quarter of the ASes blackhole more than 28 IPv4 prefixes. Thus, the overall number of blackholed prefixes is by a few members that appear to heavily rely on the RTBH service as the median value per ASN for IPv4 (IPv6) is 5 (1) blackholed prefixes. 75% of unique blackholing ASes seen across all collectors trigger blackhole prefixes at only one location. We observe that some ASes trigger blackholing at numerous locations. For instance, more than 10% of unique ASes were monitored across all collectors’ blackhole prefixes at six different locations, usually including DE-CIX Frankfurt and Giganet Kiev. Our results confirm that in practice network operators rely on blackholing services for network mitigation attacks. Therefore, understanding the IXPs’ practices for these services when coupled with RPKI is vital.

3.3 Synchronized blackholed IP prefixes

Having seen that network operators mostly blackhole very specific IPv4 address blocks (i.e., /32s), we further analyze whether these /32s share the same /24 prefix. One-third of the blackholed /32s are isolated as they solely map to different /24 address blocks. However, we find cases where the blackholed /32s cover most of their corresponding /24 prefix. Moreover, these blackholing activities are also synchronized in time meaning that the /32 prefixes belonging to the same /24 prefix have overlapping blackholing time periods. Half of all blackholed /32s are reported at the same time as related /32s (i.e. in the same /24). Figure 2 shows the number of blackholed /32s that are in the same /24 and reported at the same time. Note that we filtered out collectors with less than 150 blackholed /32s to improve the readability of the figure. Thus our plot shows only nine of the overall collectors. For almost all collectors the median value is equal to 2, meaning that we usually observe /32s blackholed by pair. Since the third quartile is below 10



(a) Average blackholing periods per prefix (b) Blackholing periods count per prefix

Fig. 3: Distribution of the (a) average blackholing periods and (b) number of periods per prefix.

/32s for most collectors, we usually observe a limited number of related /32s reported at the same time but we found a few cases where more than 100 /32s from the same /24 are blackholed at the same time (see outliers in Figure 2). We further analyze the length of time during which pairs of prefixes are blackholed together and find that one-quarter of the pair prefixes are blackholed together for less than 17 minutes. At the same time, half of the prefix pairs are blackholed for at least 1007 minutes. Thus, our results show variability in the synchronized blackholed period.

3.4 Blackholing time periods

To better understand the temporal dynamics of blackholing activities we investigate the duration and recurrence of the identified blackholed prefixes. Figure 3 shows the distribution of the average period per prefix as well as how many times a prefix is blackholed. Most of the prefixes are blackholed at most ten times, with half of the IPv4 prefixes blackholed just once. A small number of IPv4 prefixes appear to be blackholed more than 100 times and two IPv6 prefixes are blackholed 17 times.

The overall blackholing period for each prefix is half of the prefixes approximately one hour. Still, we see a lot of diversity in the blackholing periods. One-quarter of the IPv4 prefix blackholed periods last at most 10 minutes, while 10% of the blackholed periods last more than 80 hours.

4 RPKI deployment

Using data from RIPE NCC’s RPKI repository [11] we analyze the RPKI practices for blackholed prefixes. We thus fetch ROA records over the span of nine months, i.e., from January to September 2022. Our collected data comprises 336,338 IPv4 and 84,860 IPv6 RPKI ROA records registered for 34,212 ASes. For approximately 42% of ASes that have registered ROAs, we find records for





Fig. 4: Heatmap of the number of ROAs for IPv4 prefixes per prefix length and maxLength.

both IPv4 and IPv6 prefixes, while 52% and 6% of the ASes are registered only for IPv4 and IPv6 prefixes, respectively. We note that this result is consistent with the size of the IPv4 and IPv6 routing table.


Analyzing the ROAs across the five regions shows that 47%, 25% and 18% of these are registered within RIPE, APNIC and ARIN, respectively. ROAs registered with LACNIC and AFRINIC account for 8% and 1.5% of the ROAs. The observed ROA distribution is similar also for the ASes across the five registries. Through their policies, RIRs decide how long a ROA is valid and varies depending on the region. Specifically, ROAs for IPv4 prefixes registered in ARIN and AFRINIC are valid on average at least twice as long.

Breaking down the overall number of ROAs per prefix length shows that IPv4 /24 prefixes and IPv6 /48 prefixes represent approximately half of the IPv4 and IPv6 ROA records, respectively. Taking one step further we split these numbers per maxLength. Figure 4 shows the heatmap for the number of IPv4 ROAs per prefix length and maxLength. Our analysis shows that most ROAs maxLength follows the prefix length; 82.58% (277,766) of the ROAs for IPv4 prefixes satisfy this criterion. Our analysis of the IPv6 ROAs reveals similar results, i.e., for 86.75% (73,622) ROAs the maxLength is equal to the prefix length.

To ease their RPKI management organizations can register ROAs with the maximum maxLength, i.e., /32 and /128 for IPv4 and IPv6 prefixes, regardless of their prefix length. Hence a single ROA could authorize announcements for a prefix and its sub-prefixes. However, this practice is discouraged as attackers can hijack the registered resources by announcing the sub-prefixes [12]. In our dataset, we found numerous ROAs where the maximum *maxLength* is used (i.e. /32 for IPv4 and /128 for IPv6). Specifically, we find 4985 IPv4 and 592 IPv6 such ROAs which involve 942 ASes. Figure 5 plots the distribution of these ROAs per prefix length. We observe high variability among both the IPv4 and IPv6



(a) ROAs for IPv4 prefixes



(b) ROAs for IPv6 prefixes

Fig. 5: Number of ROAs with *maxLength* is equal to (a) /32 and (b) /128.

prefixes and, although discouraged, the vast majority of these IPv4 ROAs are for prefixes that are not /32s.

Focusing on our selected set of IXP members, we cross-check the advertised IP address space with the members' ROA information and filter the ROAs registered by the members. This step yields ROAs for 33,687 IPv4 and 4470 IPv6 prefixes registered for 1461 and 1031 ASNs. Among these ROAs, we also found cases where the IPv4 (IPv6) prefix length is also /32 (/128) and further focused our analysis on understanding if these ROAs are related to RTBH activities.

5 When RPKI meets RTBH

5.1 Operators RPKI management

RPKI object management can influence the effectiveness of RTBH announcements. For example, an operator can make sure that triggered /32 IPv4 or /128 IPv6 blackhole announcements are accepted by networks implementing ROV by registering the prefix in RPKI with the appropriate *maxLength* attribute beforehand. This is challenging in practice because of the fast response needed to mitigate DDoS attacks and the slow propagation of RPKI objects [7]. To further understand RPKI practice for blackholed prefixes, we devise three profiles that reflect how operators maintain their ROAs:

1. *RPKI-strict*: To comply entirely with the RPKI ecosystem operators may create very specific ROAs (/32 prefix) for IP addresses that are susceptible to DoS attacks. Thus, BGP announcements for RTBH will be RPKI valid and will be accepted by ROV-enabled networks. We believe that this approach comes with increased overhead as it requires a lot of planning that may not be possible for resources that may serve via arbitrary IP addresses.
2. *RPKI-loose*: An easier way to comply with RPKI is to create ROAs with a maxLength attribute set to /32 (/128) for IPv4 (IPv6) prefixes allowing BGP announcements for any prefix size. This practice is however strongly discouraged by the community, as it makes prefixes vulnerable to hijacks as shown by previous studies [12].
3. *RTBH-agnostic*: Operators may decide to manage their RPKI data as if they are not going to announce very specific prefixes for RTBH. IXPs’ route servers implement mechanisms to handle these cases. If an RTBH announcement is made for a prefix covered by a ROA then the route server only validates the origin ASN disregarding the ROA’s max length attribute. However, ROV-enabled networks will most likely drop these announcements.

Operator Profile	IPv4 ASes (Prefixes)	IPv6 ASes (Prefixes)
RPKI-strict	4 (4)	1 (1)
RPKI-loose	92 (553)	26 (44)
RTBH-agnostic	1438 (33177)	1021 (4731)
All Members utilizing RPKI	1461 (33687)	1031 (4770)

Table 1: Number of IXP members per network operator profile that deploy RPKI.

Using the members’ ROAs from Section 4 we identify how these networks map to the three above profiles. Note that we deliberately chose different spanning periods for the routing and RPKI data, as we investigated whether implementing RPKI comes after a series of RTBH periods. Our initial results were however inconclusive. Out of all members for the selected IXPs, 98.42% of the members are RTBH-agnostic. We find that 4 (1) members have ROAs corresponding to the IPv4 (IPv6) RPKI-strict profile and 92 (26) members have IPv4 (IPv6) ROAs with maxLength set to /32 (/128). Tables 1 list the detailed number of ASes and prefixes for each profile. Analyzing whether operators rely only on one or multiple approaches to maintain their ROAs shows that in fact most of the operators map to one single profile. We list in table 2 the number of members that map to just one single profile as well as multiple profiles. Most of the RTBH-agnostic operators choose only this approach for their ROA management. Some of these networks, however, also implement RPKI-loose or RPKI-strict practices for some of their IP addresses. Surprisingly, we find that a few operators rely only on the RPKI-loose approach.

Operator Profile	IPv4	IPv6
only RPKI-strict	0	0
only RPKI-loose	23	10
only RTBH-agnostic	1367	1004
RPKI-strict & RPKI-loose	0	0
RPKI-strict & RTBH-agnostic	2	1
RPKI-loose & RTBH-agnostic	67	16
All profiles	2	0


Table 2: Number of IXP members per *unique* network operator profile that deploy RPKI.

Fig. 6: Percentage of IXP members that blackhole IP addresses and deploy RPKI.

Operator Profile	IPv4	IPv6
only RPKI-strict	2	1
only RPKI-loose	16	1
only RTBH-agnostic	172	6
RPKI-strict & RPKI-loose	1	0
RPKI-strict & RTBH-agnostic	1	0
RPKI-loose & RTBH-agnostic	9	0

Table 3: Classification of IXP members per profile operator that *both* blackhole IP addresses and deploy RPKI.

5.2 IXP member classification

Our next step is to investigate the mapping of the devised profiles to IXP members which rely on blackholing services. To this end, for each AS we retrieve blackholed prefixes and their covering ROAs. One-third of the IXP members do not implement RPKI and are not using RTBH services. However, most of the AS that trigger blackholes also register ROAs in RPKI. Specifically, we find that 201 (8) of the 221 (15) AS that blackhole IPv4 (IPv6) prefixes also register ROAs. At the same time, our analysis reveals that few operators implement RPKI-strict and RPKI-loose practices. Table 3 provides the breakdown per profile. Unsurprisingly, we find that most of these are RTBH-agnostic. Recall that initially most of the network operators were mapped to this class. Still, our classification maps some of the operators to the other two profiles.

Figure 6 shows the percentage of ASes that trigger blackholing per operator profile. All the ASes that match the RPKI strict profile are using the RTBH services. The presence of ROA for /32 prefixes seems thus to be a good indicator of resources targeted by DDoS attacks. Unsurprisingly, blackholing ASes account for only 12.6% of all ASes that match the RTBH-agnostic profile. We find that a larger fraction (28.2%) of the RPKI-loose operators are using blackholing services for their IPv4 prefixes. Furthermore, we check whether the operators map to multiple profiles and find that 94% of the RTBH-agnostic operators are mapped only to this profile, while no operator is only RPKI-strict. Surprisingly, our




Fig. 7: Percentage of blackholed IPv4 prefixes (RTBH) and ROA prefixes (RPKI) that overlap per network per operator profile.

results show that 16 operators are only RPKI-loose, which means that these operators have poor RPKI management for all their resources.

For the ASes that use both RPKI and RTBH, we investigate the overlap between the blackholed IP space and the IP space covered by ROAs. We find that for 172 of the 201 operators the registered ROAs overlap with the blackholed IPv4 addresses, and for IPv6 all eight members have overlapping ROA and RTBH spaces. Next, we group the overlapping percentage values into four intervals, 0% - 25%, 25% - 50%, 50% - 75%, and 100%. Figure 7 shows the number of operators for each interval for the IPv4 address space. The pattern highlights the RPKI deployment. Our results show that 100 operators have registered ROAs for their entire blackholed IPv4 addresses. At the same time, we also find that 20 operators have only up to 25% of these IPs covered by ROAs. Our analysis of the members' RPKI deployment shows that for 74 of these organizations, the covering ROAs represent just up to 25% of their registered resources. For 22 members, the blackholed IPs match all their ROA records.

In addition, we analyze the time gap between the start of the blackholing period and the RPKI valid period. Most ASes triggering blackholes were already deploying RPKI for approximately one year on average. For only 120 blackholed IPv4 prefixes we find that the corresponding ROAs are registered after an average of 46 days. Our findings indicate that operators are more likely to blackhole their IP addresses whenever they suffer DDoS attacks. This in turn means that organizations place their trust in the IXPs.

6 Conclusion

RTBH and RPKI are frameworks currently used by a large number of network operators as mitigation techniques [14, 39]. In this study, we seek to understand whether and how operators combine their functionalities. Currently, there is no standard that explains how to handle both RTBH and RPKI.

Using publicly available data we first characterize separately the RTBH usage by network operators at different IXPs. Then we rely on RPKI ROA data

to devise three different operator profiles that characterize the operators' RPKI practices for RTBH announcements. RPKI-strict is impeding RTBH volatility since registering new ROAs takes significantly more time [7] compared to fast DoS response mechanisms. RPKI-loose is discouraged [12] although our results show that operators still implement this practice for both IPv4 and IPv6 prefixes. RTBH announcements from RTBH-agnostic operators are rejected by IXP members implementing ROV. Indeed ROV enabled networks disregard these RTBH announcements because they are RPKI invalid thus impeding the effectiveness of RTBH. Unless new mechanisms are designed to handle these cases, the effectiveness of RTBH is set to decrease as ROV gets more broadly deployed. Moreover, the RTBH-agnostic operators place their trust in the IXP blackholing services which implicitly increases the critical role played by IXP in Internet infrastructure.

Our study, however, comes with some limitations. Specifically, our initial set of IXP route collectors is relatively low since we study only IXPs that offer blackholing services. Also, our survey sent to the selected IXP regarding the RTBH and ROV practices only received one reply. The route collector selection process also impacts the geographic spread of the IXPs as most of the considered route collectors are based in Europe and US. Our analysis focuses only on the IXP members which represent a small fraction of the overall number of ASes.

Our results show that RTBH agnostic is the most common practice, so quantifying the current impact of ROV-enabled networks on RTBH is a good follow-up to this work. That is however out of scope for this paper as it requires significant research efforts in surveying which IXP members are implementing ROV and monitoring the way they handle RTBH announcements.

The considered three RPKI practices, however, do not offer a perfect solution in combining the two security frameworks. Thus, our current recommendation is that network operators follow current best practice of implementing a *RTBH-agnostic* RPKI approach, leaving to the IXP's the duty of checking the validity of RTBH requests. However, in the long term, operators and researchers should make an effort in devising and standardizing a technique that unambiguously describes the coexistence of RPKI and RTBH.

Acknowledgments

We thank the anonymous reviewers and our shepherd for their helpful comments. This research was supported in part by the MANRS Fellowship Program. Ioana Livadariu was partially funded by the SimulaMet's internal funding. We would like to thank the IXPs who responded to our survey.

A Ethics

We are not aware of any ethical issue raised by this work. Our analysis relies on publicly available datasets.

References

1. Daniel Wagner, Daniel Kopp, Matthias Wichtlhuber, Christoph Dietzel, Oliver Hohlfeld, Georgios Smaragdakis, and Anja Feldmann. United we stand: Collaborative detection and mitigation of amplification ddos attacks at scale. In *Proceedings of the 2021 ACM SIGSAC conference on computer and communications security*, pages 970–987, 2021.
2. Jakub Czyz, Michael Kallitsis, Manaf Gharaibeh, Christos Papadopoulos, Michael Bailey, and Manish Karir. Taming the 800 pound gorilla: The rise and decline of ntp ddos attacks. In *Proceedings of the 2014 Conference on Internet Measurement Conference, IMC '14*, page 435–448, New York, NY, USA, 2014. Association for Computing Machinery.
3. Christoph Dietzel, Anja Feldmann, and Thomas King. Blackholing at ixps: On the effectiveness of ddos mitigation in the wild. volume 9631, pages 319–332, 03 2016.
4. Daniel Kopp, Christoph Dietzel, and Oliver Hohlfeld. Ddos never dies? an ixp perspective on ddos amplification attacks. In *Passive and Active Measurement*. Springer International Publishing, 2021.
5. Thomas King, Christoph Dietzel, Job Snijders, Gert Döring, and Greg Hankins. BLACKHOLE Community. RFC 7999, Oct 2016.
6. M. Lepinski and S. Kent. An Infrastructure to Support Secure Internet Routing. <https://datatracker.ietf.org/doc/html/rfc6480>, 2012.
7. R. Fontugne, A. Phokeer, C. Pelsser, K. Vermeulen, and R. Bush. RPKI Time-of-Flight: Tracking Delays in the Management, Control, and Data Planes. In *Proceedings of PAM'23*. LNCS, 2023.
8. Khwaja Zubair Sediqi, Lars Prehn, and Oliver Gasser. Hyper-specific prefixes: gotta enjoy the little things in interdomain routing. *ACM SIGCOMM Computer Communication Review*, 52(2):20–34, 2022.
9. Y Gilad, S Goldberg, K Sriram, J Snijders, and B Maddison. Rfc 9319 the use of maxlen in the resource public key infrastructure (rpki). 2022.
10. Packet Clearing House. Internet Exchange Directory. <https://www.pch.net/ixp/dir>, 2023.
11. RIPE NCC. RPKI repository archive. <https://ftp.ripe.net/rpki/>, 2023.
12. Gilad Yossi, Sagga Omar, and Goldberg Sharon. Maxlength considered harmful to the rpki. CoNEXT '17, New York, NY, USA, 2017. Association for Computing Machinery.
13. C. Lynn, S. Kent, and K. Seo. X.509 Extensions for IP Addresses and AS Identifiers. <https://www.rfc-editor.org/rfc/rfc3779>, 2004.
14. DE-CIX. RPKI at the DE-CIX route servers. <https://www.de-cix.net/en/resources/service-information/route-server-guides/rpki>, 2023.
15. Equinix. Resource Public Key Infrastructure (RPKI). <https://docs.equinix.com/en-us/Content/Interconnection/IX/IX-rpki.htm>, 2023.
16. The BIRD Internet Routing Daemon. <https://bird.network.cz/>.
17. Flavio Luciani. Checking prefix filtering in IXPs with BIRD and OpenBGPD. <https://blog.apnic.net/2021/11/15/checking-prefix-filtering-in-ixps-with-bird-and-openbgpd/>, 2023.
18. FranceIX. RAPPORT TECHNIQUE Q1 2020. <https://blog.franceix.net/rapport-technique-q1-2020/>, 2020.
19. PeeringDB: T-CIX Route Servers. <https://www.peeringdb.com/net/8295>.
20. Diego Neto (NL-ix). BIRD route-server configuration: click, done! <https://indico.uknof.org.uk/event/39/>, 2017.

21. Vasileios Giotsas, Georgios Smaragdakis, Christoph Dietzel, Philipp Richter, Anja Feldmann, and Arthur Berger. Inferring bgp blackholing activity in the internet. In *Proceedings of the 2017 Internet Measurement Conference*, IMC '17, page 1–14, New York, NY, USA, 2017. Association for Computing Machinery.
22. Christoph Dietzel, Matthias Wichtlhuber, Georgios Smaragdakis, and Anja Feldmann. Stellar: Network Attack Mitigation Using Advanced Blackholing. In *Proceedings of the 14th International Conference on Emerging Networking EXperiments and Technologies*, CoNEXT '18, page 152–164, New York, NY, USA, 2018. Association for Computing Machinery.
23. Loïc Miller and Cristel Pelsser. A taxonomy of attacks using bgp blackholing. In *Computer Security – ESORICS 2019*. Springer International Publishing, 2019.
24. Matthias Wichtlhuber, Eric Strehle, Daniel Kopp, Lars Prepens, Stefan Stegmüller, Alina Rubina, Christoph Dietzel, and Oliver Hohlfeld. Ixp scrubber: Learning from blackholing traffic for ml-driven ddos detection at scale. In *Proceedings of the ACM SIGCOMM 2022 Conference*, SIGCOMM '22, page 707–722, New York, NY, USA, 2022. Association for Computing Machinery.
25. Florian Streibelt, Franziska Lichtblau, Robert Beverly, Anja Feldmann, Cristel Pelsser, Georgios Smaragdakis, and Randy Bush. BGP Communities: Even more Worms in the Routing Can. In *Proceedings of ACM IMC 2018*, Boston, MA, October 2018.
26. Andreas Reuter, Randy Bush, Italo Cunha, Ethan Katz-Bassett, Thomas C. Schmidt, and Matthias Wählisch. Towards a rigorous methodology for measuring adoption of rpki route validation and filtering. 48(1):19–27, apr 2018.
27. Taejoong Chung, Emile Aben, Tim Bruijnzeels, Balakrishnan Chandrasekaran, David Choffnes, Dave Levin, Bruce M. Maggs, Alan Mislove, Roland van Rijswijk-Deij, John Rula, and Nick Sullivan. Rpki is coming of age: A longitudinal study of rpki deployment and invalid route origins. In *Proceedings of the Internet Measurement Conference*, IMC '19, page 406–419, New York, NY, USA, 2019. Association for Computing Machinery.
28. Job Snijders, Mikael Abrahamsson, and Ben Maddison. Resource Public Key Infrastructure (RPKI) object profile for Discard Origin Authorizations (DOA). Internet-Draft draft-spaghetti-sidrops-rpki-doa-00, Internet Engineering Task Force, Mar 2022. Work in Progress.
29. Packet Clearing House. Pch raw routing data. https://www.pch.net/resources/Raw_Routing_Data/. (Accessed on 05/25/2023).
30. DE-CIX. Blackholing guide. <https://www.de-cix.net/en/resources/service-information/blackholing-guide>, 2023.
31. Equinix. Remotely Triggered Black Hole. <https://docs.equinix.com/en-us/Content/Interconnection/IX/IX-rtbh-guide.htm>, 2023.
32. Giganet. Blackhole (BGP). <https://giganet.ua/en/service/blackhole>, 2023.
33. FranceIX. Blackholing. <https://www.franceix.net/fr/services/infrastructure/blackholing>, 2023.
34. SeattleIX. Blackholing. <https://www.seattleix.net/blackholing>, 2023.
35. Techtargget Security. Major ddos attacks increasing after invasion of ukraine. <https://www.techtargget.com/searchsecurity/news/252521150/Major-DDoS-attacks-increasing-after-invasion-of-Ukraine>, June 2022. (Accessed on 05/25/2023).
36. The Record. Ddos attacks surge in popularity in ukraine — but are they more than a cheap thrill? <https://therecord.media/ddos-attacks-surge-in-popularity-in-ukraine-but-are-they-more-than-a-cheap-thrill>, July 2022. (Accessed on 05/25/2023).

37. Computer Weekly. Ukraine war drives ddos attack volumes ever higher. <https://www.computerweekly.com/news/252523959/Ukraine-war-drives-DDoS-attacks-volumes-ever-higher>, August 2022. (Accessed on 05/25/2023).
38. National Cyber Security Center. Uk government assess russian involvement in ddos attacks on ukraine. <https://www.ncsc.gov.uk/news/russia-ddos-involvement-in-ukraine>, February 2022. (Accessed on 05/25/2023).
39. National Institute of Standards and Technology (NIST). RPKI-ROV History of Unique Prefix-Origin Pairs (IPv4). <https://rpki-monitor.antd.nist.gov/ROV>, 2024.