http://researchcommons.waikato.ac.nz/

Research Commons at the University of Waikato

Copyright Statement:

The digital copy of this thesis is protected by the Copyright Act 1994 (New Zealand).

The thesis may be consulted by you, provided you comply with the provisions of the Act and the following conditions of use:

- Any use you make of these documents or images must be for research or private study purposes only, and you may not make them available to any other person.
- Authors control the copyright of their thesis. You will recognise the author's right to be identified as the author of the thesis, and due acknowledgement will be made to the author where appropriate.
- You will obtain the author's permission before publishing any material from the thesis.

Measuring the Effectiveness of Routing Defenses Through the Lens of DROP

A thesis

submitted in partial fulfilment

of the requirements for the Degree

of

Master of Cyber Security

at

The University of Waikato

by

Leo Oliver-Dowling

THE UNIVERSITY OF
WAIKATO
Te Whare Wānanga o Waikato

2022

Abstract

This work analyzes the properties of 712 prefixes that appeared in Spamhaus' "Don't Route Or Peer" (DROP) list over a nearly three-year period from June 2019 to March 2022. The 712 known abused prefixes are used as a lens to assess the current threat landscape and evaluate several of the leading routing defense mechanisms. A thorough characterization of these 712 prefixes is performed and it is found that a larger fraction of the hijacked prefixes were from Regional Internet Registries (RIRs) with restrictive policies regarding Resource Public Key Infrastructure (RPKI) eligibility. It is also found that attackers were predominately targeting address space that was unrouted and not RPKI-signed. The work reveals that attackers were subverting multiple defenses against malicious use of address space, including creating fraudulent Internet Routing Registry records for prefixes shortly before using them. Other attackers disguised their activities by announcing routes with origin Autonomous Systems (ASes) consistent with historic route announcements, and in one case, with the Autonomous System Number (ASN) in a RPKI Route Origin Authorization. Finally, the work quantifies the substantial and actively-exploited surface in unrouted space, which warrants reconsideration of RPKI eligibility and policies by both operators and RIRs.

Acknowledgements

I would like to thank my supervisor, Associate Professor Matthew Luckie, for your patience, guidance, and continuous support throughout this project. I have benefited greatly from your wealth of knowledge and I am extremely grateful that you took me on as a student.

I would like to thank the Gallagher Group for awarding me the Sir William Gallagher Cyber Security Scholarship. Without receiving this financial support completing this research would not have been possible.

Finally, I would like to thank my family and friends, specifically my partner for supporting throughout the past year.

Contents

Chapter 1

Introduction

1.1 The Problem

Malicious use of Internet address space has been a persistent threat for decades. In some cases this malicious use involves an actor falsely asserting ownership of addresses they do not in fact own. In other cases a malicious actor uses its own address space for fraudulent activity such as spam or malware distribution. They may obtain such addresses fraudulently, e.g., by forging documentation needed to procure it, or they may acquire it from hosting companies that knowingly lease address space for malicious use. Malicious actors enjoy advantages against those trying to thwart their activities, the biggest advantage being the difficulty of detecting illegitimate users of address space in real time, especially at high traffic rates common in busy parts of the Internet. Compounding the challenge, privacy concerns have triggered legislation [47, 16] that has reduced accessibility to data that enables attribution of Internet address and domain name ownership [1]. Privacy concerns have also led to development and deployment of encrypted protocols [31], increasing the challenge of network defenders more generally.

There have been at least four classes of approaches to prevent and detect address space abuse. The first class is the maintenance of databases of legitimate address ownership, from the myriad Internet Routing Registry (IRR)

databases that launched in the 1990s [28] to today's cryptographically signed Regional Internet Registry (RIR) based Resource Public Key Infrastructure (RPKI) [24]. An IRR is a database used to share routing information, most importantly it allows an operator to share which Autonomous System(s) (AS(es)) are authorized to originate a prefix. RPKI also proves the association between specific prefixes and their holders, however, it does so in a manner that can be cryptographically validated. In both IRR, and RPKI it is up to operators to decide whether they filter based on the invalidity of a route. The IRR databases lack validation for who can register routing information and they have proved hard to accurately maintain. RPKI still faces obstacles, operators lack incentive to deploy it, misconfigurations are common leading to operators reluctance to filter based on its invalidity, and it does not protect against hijacks where the hijacker manipulates an AS path and adds the AS authorized to originate the prefix to the end of the path. The second class of approach is a more thorough authentication of the entire routing announcement, not just the originating network, e.g. BGPSec [25, 9]. Efforts in this direction have continued for decades however they all require modifications to the Border Gateway Protocol (BGP) protocol, which is an obstacle to deployment given the vast installed base of BGP software and devices, so actual deployment may take at least one more decade. The third class of defense is automated detection of illicit use of address space once it has started; route hijack detection [29, 30, 35, 64, 57]. Detection techniques have two inherent limitations: (1) they are reactive – they respond to hijacks after they occur, and (2) it is challenging to discern hijacks from many legitimate routing changes, and so detection approaches suffer from false positives [43]. Finally, the fourth class of defense is the use of blocklists, which many companies distill from third-party reporting and sell to network defenders [38]. These blocklists consist of IP addresses or domain names that are implicated in malicious activity, which operators then use to drop traffic to/from. Despite many attempts to detect and mitigate BGP hijacks, they are a persistent problem [46, 7, 21, 58, 15, 14].

1.2 Contributions of This Work

Through the lens of one of the most respected blocklists on the Internet: Spamhaus Don't Route Or Peer (DROP) list [37] we undertook a study of the effectiveness of IRR, RPKI and the DROP list itself. Spamhaus investigators regularly update the DROP list with IPv4 address prefixes that pose a presumed threat to the Internet community[60]. We use the prefixes added to the DROP list over a nearly three-year period from June 2019 to March 2022, noting the prefixes that appear and the date that they appeared. We use DROP for four reasons. First, it is well documented – each entry describes why it was added [63]. Second, it represents the most seriously abused prefixes, which Spamhaus encourages operators to refuse traffic to/from. Third, a human validates the decision to add prefixes to the DROP list, increasing its accuracy [61]. Fourth, access to this data is free, enabling others to more easily reproduce this work.

The findings and contributions of this work are as follows. We characterize the address space that appeared in DROP over the nearly 3 year period from June 5, 2019 to March 30, 2022. Through this characterization it is found that North America and Africa have larger portions of legacy address space (large blocks of address space allocated before the establishment of the RIR system in December, 1997) in DROP labeled *hijacked*, which may be a direct result of their policy that prohibits legacy resource holders from RPKI-signing their addresses. The work reveals that attackers are circumventing defenses against malicious use of address space, including (1) registering IRR records for prefixes shortly before the attacker uses those prefixes, (2) announcing routes with origin ASes consistent with historic route announcements, and (3) announcing routes with the RPKI-signed origin. Encouragingly, the process of an owner reclaiming their prefix and having it removed from DROP spurred RPKI adoption: prefixes removed from DROP were RPKI-signed at a higher rate (42.5%) than prefixes that were not added to DROP (22.3%). However, current Regional Internet Registry (RIR) policy around RPKI signing unallo-

cated address space provides a vulnerability that attackers are exploiting, with 40 unallocated prefixes appearing in DROP during our study period. Further, of the 36.7 /8 equivalents of allocated but unrouted address space, 6.7 (18.3%) were RPKI signed in a way that would allow an attacker to hijack the address space. Because deployment of AS path authentication mechanisms [25, 9] may take at least one more decade, our analysis demonstrates the interim benefit of AS0 policies to enable operators to protect their own address space from abuse.

1.3 Overview

Chapter 2 provides the reader background on the incentive for this work. Firstly, it covers the foundations of BGP and more specifically how *BGP hijacking* is possible. Secondly, the two most prominent defense mechanisms against *BGP hijacks* which will be analyzed in this work (RPKI, and IRR) are explained, along with their inherent limitations, and policies that have impacted their effectiveness. In Chapter 3, for each of the public datasets that are used in this analysis there is a description that covers, what the dataset is comprised of, how it was used in this analysis, and where it was obtained. Chapter 4 performs a thorough characterization of the 712 prefixes that appeared in DROP. Each section characterizes the prefixes based on a different quality: §4.1 classifies the prefixes based on Spamhaus's own description of the prefixes activity, §4.2 classifies the prefixes based on their allocating RIR, §4.3 classifies them based on their allocation status at the time they appeared in DROP, and finally §4.4 attempts to infer their activity based on features observed through BGP. Each section provides greater understanding of the types of abuse that the Internet is currently facing, and discusses the implications of this abuse. Chapter 5 examines the effect that a prefix being added to DROP had on the routing system, and on the behavior of both attackers and network operators. Chapters 6 and 7 examine the effectiveness of IRR and

RPKI respectively. They look at what fraction of the DROP prefixes had these two mechanisms deployed and for the prefixes that did have them deployed, they analyze why the mechanisms were ineffective at preventing the prefixes from being used for abuse. Finally in Chapter 8 the study is concluded with a discussion of the implications of this works findings and suggestions for future solutions.

Chapter 2

Background & Related Work

This chapter first provides an explanation of BGP and describes the underlying vulnerability in BGP that makes BGP hijacking possible. It then looks at previous work that has examined how malicious actors are exploiting this vulnerability in BGP. The two address use authorization systems that we evaluate in this work, IRR, and RPKI are then explained, along with their limitations. Lastly, policies held by RIRs that have significantly inhibited the deployment of the RPKI are mentioned.

2.1 BGP Hijacks

The Border Gateway Protocol (BGP) is a fundamental component of networking which is responsible for communicating routing policies between distinct networks known as Autonomous Systems (ASes), creating the mesh of interconnected networks which we know as the Internet. An AS is a large network or group of networks that has a unified routing policy, each AS has a unique identifier known as an ASN. An AS uses BGP to announce a block of IP addresses (these blocks of IP addresses are referred to as prefixes) for which it has routes, the AS that makes this announcement is the *origin AS*. These announcements are then propagated to neighboring ASes creating a path of ASNs called the *AS path*, which can be used to reach that prefix. Because BGP relies on trust between ASes, an AS can announce a prefix which it does not own.

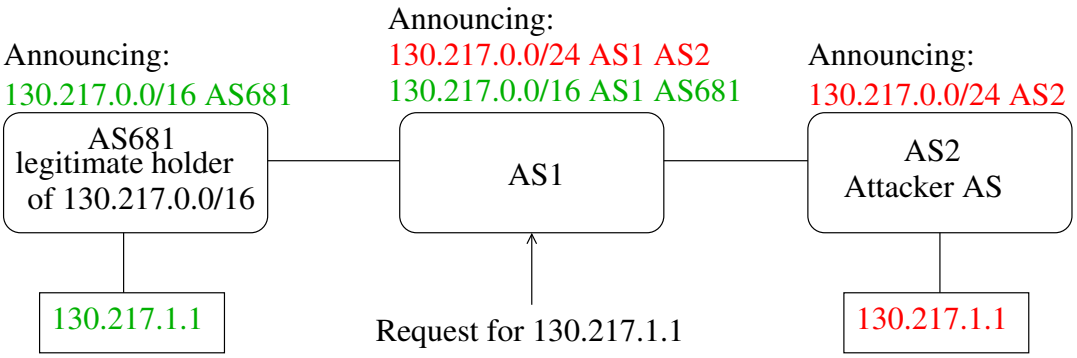

Figure 2.1: An example of a BGP hijack. The University of Waikato AS681 legitimately announces their /16. The attacker AS2 announces a more specific /24 within Waikatos /16. When a request is made to AS1 for the address 130.217.1.1 (which falls within both the /16 and /24), BGP favors the more specific route i.e. the /24, and so traffic is directed to the attacker AS2.

When an AS does so without permission (illegitimately), it is a *BGP hijack*. Figure 2.1 shows a theoretical example of an attacker (AS2) hijacking a more specific prefix within the University of Waikato prefix 130.217.0.0/16. In 2002, accidental misconfigurations were a common type of hijack [40]. Intentional hijacks have been used by cyber criminals for interception of traffic destined to the hijacked addresses [10], or as a way to use the addresses themselves to send spam [51] or perform large-scale DDoS attacks [66].

In 2015, Vervier et al. [72] illuminated the presence of hijacks occurring regularly in the wild. They manually examined the BGP behavior of 437 prefixes that sent spam to spam traps under their control, and found 64 had behavior that resembled a hijack. All 437 prefixes were unannounced by their owner prior to being hijacked. For 5 of these 64 prefixes, the hijacker forged the AS path by using the ASN that previously originated the prefix as the origin of their paths, making the announcement falsely appear as if originated by the legitimate owner [49, 23].

In 2018 Testart et al. [65] profiled the behavior of serial hijacker ASes that repeatedly offend by hijacking lot of different prefixes. The authors acquired a set of known hijacking ASes from network operator mailing lists to analyze the characteristics and behavior patterns of these hijacking ASes and how they

differed from legitimate ASes. They used this knowledge to train a classifier to infer other ASes with similar features to be hijackers. Of the 19,103 ASes in their prediction set, their classifier found 934 ASes having similar behavior to serial hijackers.

2.2 Address use authorization systems

As early as the 1990s, recognition of the need to provide some accountability and attribution around IP address work led to the creation of systems to allow ISPs to register information regarding how they will use their resources, so that ISPs can make appropriate routing decisions.

2.2.1 IRR

The Internet Routing Registry (IRR) is a set of 25+ distributed routing information databases that use the Routing Policy Specification Language (RPSL) to enable network operators to share information [28]. The Routing Assets Database (RADb) – operated by Merit – is the most complete IRR, as in addition to being an IRR itself, it mirrors various other IRRs [42]. The most important record in the IRR is the *route object*, which contains the IP prefix and origin AS that a network intends to announce in BGP.

Lack of incentive for network operators to maintain accurate records in the IRR has reduced its utility. Worse, lack of validation of registration data renders the IRR vulnerable to abuse by attackers who can easily register false information and even revoke registrations of others' address space [55, 45, 20].

2.2.2 Resource Public Key Infrastructure

The integrity limitations of the IRRs ultimately led to development of the now recommended approach to prevent unauthorized use of address space: the Resource Public Key Infrastructure (RPKI) [36, 24]. The RPKI supports cryptographic attestation that a network, identified by their ASN, is authorized

ROAs:

Origin	Prefix	Max Length
AS681	130.217.0.0/16	/16
AS682	130.217.0.0/16	/16

Routes:

Origin	Prefix	Validity	
AS681	130.217.0.0/16	Valid	
AS682	130.217.0.0/16	Valid	
AS2	130.217.0.0/16	Invalid	← Invalid Origin
AS3	130.217.0.0/24	Invalid	← Invalid Origin & Prefix Length

Figure 2.2: An example of two ROAs covering a single prefix. An RPKI route is valid if ANY of the ROAs are valid therefore both AS681 and AS682 can make valid announcements for 130.217.0.0/16. The bottom two routes are examples of invalid origin ASes, the bottom route additionally has an invalid prefix length as the length of the prefix is more specific than that of the ROA.

to *originate* a route for a prefix into the global routing system (known as a *route origin authorization* or ROA). Each of the five RIRs have their own key that they use in the signing of ROAs provided by their members. A ROA may contain an ASN that is permitted to originate a prefix, or AS0 if the prefix, and any more specific prefix within, should not be routed. A route is *RPKI-valid* if *any* ROA asserts the announcement as valid – i.e. the origin AS matches a ROA for the prefix, and the prefix length is less than or equal to the maximum prefix length contained in the ROA. A route is *RPKI-invalid* if there is one or more ROAs for a prefix, but the origin AS does not match the ASN in any of the available ROAs, or the ROA's maximum prefix length attribute is less than the length of the prefix contained in the route. Examples of both valid and invalid routes are shown in Figure 2.2. Both RIRs and individual networks can create AS0 ROAs – a network can assert that their allocated address space should not be routed, and an RIR can assert that unallocated address space in their free pool should not be routed.

Success of the RPKI at preventing origin hijacks requires two sides of participation: networks must register their own prefixes in the RPKI, and networks

must drop all BGP assertions for prefixes that are RPKI-invalid. This latter practice is called *route origin validation* (ROV). The use of ROV will prevent origin hijacks of prefixes that have valid ROAs.

RPKI misconfigurations may result in the flagging of legitimate announcements as invalid [73], a risk that has slowed the deployment of ROV [13]. Network operators began gradually deploying RPKI in 2011. As of March 2022, approximately 35% of observably routed prefixes have RPKI-valid announcements [44]. However, RPKI does not protect from rogue ASes from *propagating* invalid announcements from their neighbors, nor from *BGP path manipulation hijacks* [12], where a hijacker forges the legitimate owner's ASN for the origin of the prefix. Vervier et al. highlighted that hijacks using the owner's ASN were still a possibility with RPKI [72]. This work identifies a real world example of an attacker that circumvented RPKI and performed a RPKI-valid hijack.

2.3 The Regional Internet Registry System and IPv4 Exhaustion

Originally the IPv4 registry was just a list of IP address ranges and the organisations to whom they had been allocated to. As more entities joined the Internet, this list of allocations grew so the Internet Assigned Numbers Authority (IANA) was created to facilitate this responsibility. Despite this, as a result of the Internet's rapid expansion, it became evident that IANA would not be able to scale to fulfill the range of different regional requirements. And so five regional level subsidiary organisations were established to assume allocation in collaboration with IANA. The five RIRs that were established are: AFRINIC, APNIC, ARIN, LACNIC, and RIPE.

In early 2011, IANA assigned each of the 5 RIRs with the final /8 blocks of IPv4 address space. ARIN exhausted its address space by 2015. APNIC, RIPE, and LACNIC deployed policies to ease the rate of address space deple-

tion once they reached their last /8 blocks. These policies limited allocations to members to a maximum prefix size of /22. Because AFRINIC had fewer members than other RIRs, there was less demand for address space, and they were slower to reach their final /8. It was not until reaching their last /11 in 2020 that AFRINIC enforced a policy limiting the prefix allocation size to a /22 [3]. This made AFRINIC a target for fraudulent address requests [18].

2.4 Policies Inhibiting RPKI Uptake

2.4.1 RPKI Signing of Legacy Address Space

Legacy prefixes are prefixes that were directly allocated to a recipient before the inception of the RIR system in December, 1997 and therefore, their holder may not have any affiliation with any of the five RIRs. Legacy resource holders have stronger property rights than operators who obtain prefixes from RIRs [39]. This leads to a tussle between prefix holders and the RIRs, as only an RIR can RPKI-sign a prefix. ARIN and AFRINIC will not RPKI-sign legacy resources unless the resource holder signs their Registration Service Agreement (RSA) [4, 8]. Signing an RSA with an RIR causes the resource holder to lose all of their legacy rights therefore legacy holders in ARIN, and AFRINIC are often reluctant to sign these RSAs [39]. APNIC, LACNIC and RIPE will RPKI-sign legacy address space without the resource holder signing their RSA, provided the legacy resource holder can prove their ownership of the resource [56, 32, 33, 52].

Legacy address space also has a smaller fraction of used addresses when compared to allocated address space, as many organizations acquired much more IPv4 address space than they needed [11]. For example the University of Waikato acquired a block of 65,536 addresses but will never need this many, as did most other universities in New Zealand. The policy challenges in RPKI-signing legacy space [4, 8], combined with the fact that legacy address space is often left unannounced, render legacy address space an easy target for a

hijacker.

2.4.2 Using AS0 to Assert Invalidity

Because unallocated address space should not be routed, several RIRs have discussed policies to allow the RIR itself to create AS0 ROAs for unallocated address space. As of May 2022, APNIC and LACNIC have implemented such policies to create AS0 ROAs, however, they sign these ROAs with a key that is not configured by default in RPKI validation software [5, 48]. APNIC's public statement is that operators should only use these ROAs for informational purposes, and not to reject BGP announcements that defy those assertions [5]. These factors inhibit the deployment of defenses that would prevent abuse of unallocated address space.

Most operators rely on the RIRs to issue and maintain ROAs on their behalf, which operators configure through a web interface provided by the RIR. While LACNIC does create AS0 assertions for unallocated address space, its web interface does not allow operators to create AS0 ROAs for their allocated resources as of May 2022 [34]. This inhibits the deployment of defenses that prevent the abuse of allocated but unrouted address space in the LACNIC region.

Chapter 3

Data Sources

3.1 Spamhaus DROP and SBL Database

Spamhaus is an international organisation that supplies realtime threat intelligence information to the Internet's major networks. Spamhaus compiles several widely used blocklists, including the Don't Route Or Peer (DROP) list of IPv4 prefixes that Spamhaus deems pose a threat to Internet users [60]. Spamhaus manually confirms serious evidence of malicious activity before adding a prefix to the DROP list [61]. The DROP list contains a list of IPv4 prefixes that Spamhaus recommends a network not accept any packets from. Spamhaus also maintains the Spamhaus Block List (SBL) database, which records a justification for why each prefix is in the DROP list [63]. Spamhaus assigns each unique entry in the DROP list an identifier that allows operators and researchers to query the SBL database to access this information. We used daily snapshots of the DROP list that Firehol published in its compilation of threat intelligence [67] between June 5, 2019 and March 30, 2022– a nearly 3 year period. Over this time, Spamhaus added 712 unique prefixes to the DROP list. One feature of the SBL database is that prefixes that Spamhaus removes from the blocklist no longer have public SBL records. Overall, the SBL records for 526 of the total 712 prefixes that appeared in DROP (73.9%) were acquired for this work.

3.2 Supplementary Data Sets

In addition to the Spamhaus DROP blocklist data set, this work used four additional data sources to support analysis and interpretation of the blocklist data.

3.2.1 BGP Routing Data

Route Views is a project created by the University of Oregon's Advanced Network Technology Center that allows researchers and network operators to observe global BGP routing information from various perspectives of the Internet. Every two hours, the RouteViews project publishes snapshots of the routing tables provided by peers of route collectors operated by RouteViews. These routing table snapshots contain a list of the prefixes that are observed by each peer and the AS path to reach that prefix. This work analyzed the BGP announcements made near the time a prefix appeared in DROP recorded by all 36 Route Views collectors to investigate routing features that may have led Spamhaus to blocklist a prefix.

3.2.2 IANA IPv4 Address Space Registry

The Internet Assigned Numbers Authority (IANA) is a standards body that manages global IP address allocations and ASN allocations [26]. IANA maintains a list for the allocation status of all /8 blocks of IPv4 address space. This list defines whether a /8 block was delegated to an RIR, was allocated directly to an AS before the establishment of RIRs (legacy), or is reserved by IANA for a specific application or future use. IANA's list has a granularity of /8 blocks and does not account for legacy address space that is returned to an RIR and then reallocated to a different AS, address space that is delegated to an RIR but not allocated to an AS, or address space that the holder has transferred to a different RIR from the RIR managing the delegation as recorded by IANA. Since this work requires knowledge of whether a prefix has

remained legacy (did not affiliated with an RIR) in order to accurately make claims related to legacy address space, a method to determine whether a prefix is currently legacy is devised in Section 3.2.4 which utilizes the IANA Address Space Registry dataset and the WHOIS database.

3.2.3 RIR Statistics

Each of the five RIRs publish "stats" files that provide a snapshot of the status of Internet number resources [6], and the RIRs archive these files. These files adhere to the RIR statistics exchange format [6]. RIPE maintain an FTP server that contains current and historical statistics files for each of the RIRs [54]. This work used each of the stats files to find the status of each prefix at the date that the prefix appeared in DROP. The status of a prefix is either *allocated* or *assigned* meaning that a RIR has allocated the prefix to an AS, or *reserved* or *available* which is how the RIRs label address space that they have not yet allocated to an AS. We use these stats files to determine which RIR a prefix was allocated to and what the status of the prefix was at the time it appeared in DROP.

3.2.4 WHOIS Database

To access information regarding the holder of address space, RIRs offer public access registration data via their WHOIS service. The WHOIS protocol is a query and response protocol used to query the databases containing registered users or assignees of an Internet resource [17]. For DROP prefixes within an IANA legacy /8 block, we use the prefixes registration date in the WHOIS database to determine whether each prefix remains legacy or has become affiliated with a RIR. If the registration date is before the establishment of the RIR system (11 December 1997) then we infer that the prefix has retained its legacy status, and if it was after this date then we infer that the prefix has affiliated with a RIR.

route:	130.217.0.0/16
descr:	On behalf of REANNZ Member
origin:	AS681
mnt−by:	MAINT−AS38022
changed:	noc@reannz.co.nz 20140516
source:	RADB

Figure 3.1: A IRR route object for the University of Waikato prefix 130.217.0.0/16, in the RADb.

3.2.5 IRR

Merit archive daily snapshots of the RADb [50]. which this work utilizes to view historical IRR records relating to the prefixes in DROP. Figure 3.1 shows an example of the fields that comprise an IRR route object. The fields are described as follows: *route* contains the prefix of the route to be originated, *descr* contains text that describes the route, *origin* contains the AS that is authorized to originate the route, *mnt-by* contains the ORG-ID of the organization that configures (maintains) the IRR object, *changed* specifies who submitted the update and when the object was updated, and *source* contains the name of the routing registry.

3.2.6 RPKI

RIPE maintains an archive of historic RPKI ROAs for each RIR [53]. RIPE publish daily RPKI ROAs to this archive which contain lists of RPKI signed prefixes and the AS they are signed to. We parse the data from every RIRs ROA files for each day from June 5, 2019 to March 30, 2022 and determine if a DROP prefix is RPKI signed by checking it is within a prefix in a ROA file or not.

Chapter 4

Characterizing Prefixes in DROP

4.1 By Spamhaus Category

The SBL record (which is free form text) for each prefix was processed and each prefix was and placed into one or more of the following categories:

1. **Hijacked (HJ)**. Prefixes an attacker obtains through fraud from an RIR or through announcing a prefix that an RIR assigned to another network.

2. **Snowshoe Spam (SS)**. Prefixes are used by spammers to originate spam from many IP addresses within a prefix, in order to evade detection.

3. **Known Spam Operation (KS)**. Prefixes under the control of, or otherwise connected with, a spam operation.

4. **Malicious Hosting (MH)**. Prefixes used by bulletproof hosting services knowingly hosting malicious actors.

5. **Unallocated (UA)**. Prefixes that neither IANA nor an RIR has allocated to an AS (but attackers are using). These prefixes are also often referred to as 'bogon' prefixes.

Record	Keyword	Classification
SBL310721	AS204139 spammer **hosting**	*malicious hosting*
SBL240976	**hijack**ed IP range ... billing@a**hosting**inc.com	*hijack*
SBL502548	**Snowshoe** IP block on **Stolen** AS62927 ... james.johnson@networx**hosting**.com	*snowshoe*, *hijack*
SBL322513	Register Of **Known Spam Operations** ... **snowshoe** range	*known spam operation*, *snowshoe*
SBL294939	Register Of **Known Spam Operations** ... illegal netblock **hijacking** operation	*known spam operation*, *hijack*
SBL325529	Department of Defense ... Spamhaus believes that this IP address range is being used or is about to be used for the purpose of high volume spam emission.	*snowshoe*

Table 4.1: Excerpts from SBL records that I used to classify DROP prefixes.

6. **No SBL Record (NR)**. Prefixes that the SBL record could not be obtained for, because Spamhaus had removed the record after the prefix holder had remediated and before this work was started.

To categorize each prefix, each SBL record was searched for the case-insensitive strings: 'hijack' + 'stolen', 'snowshoe', 'known spam operation', 'hosting', and 'unallocated'+'bogon', as illustrated in Table 4.1. The word 'hosting' in the first record of Table 4.1 led to that SBL record being classified as *malicious hosting*. It was manually verified that Spamhaus used 'hosting' in relation to a malicious activity – e.g. spam hosting, bulletproof hosting, botnet

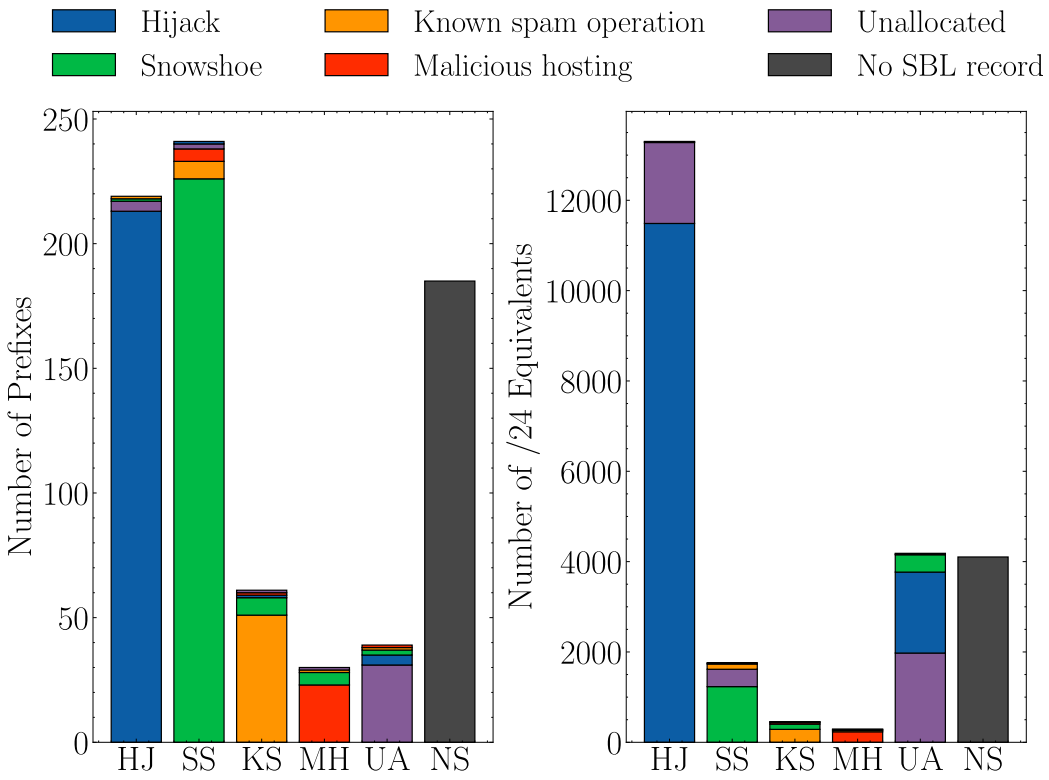

Figure 4.1: Classification of Spamhaus DROP entries. The bottom segments of the bars are portions of prefixes (left) and /24 equivalents (right) with the corresponding classification exclusively. The segments above are portions of entries with the second labeled classification additionally. Although prefixes with the label *snowshoe* take up nearly a third of prefix additions in DROP, they are small prefixes and comprise only 4.76% of address space. In contrast the *hijack* and *bogon* categories contain large prefixes and therefore have larger portions of address space.

hosting etc, to avoid spurious classifications that could occur when hosting was not used in that context, such as in the second and third records in Table 4.1. 90% of SBL records contained one keyword, 2.7% of SBL records contained two keywords, and the remaining 7.3% of SBL records contained none. For this last category the prefix's category was manually inferred, e.g., The last record in Table 4.1 (SBL325529) was classified as *snowshoe spam* because Spamhaus had reason to believe that the IP range could be used for high volume spam emission. There were two prefixes where Spamhaus did not provide enough information to infer an accurate label, so these prefixes were not included in any category reliant analysis.

The *hijack* and *snowshoe* categories were the largest, each around a third of the prefixes in DROP as shown in Figure 4.1. The categories show little overlap: *snowshoe* prefixes show the most overlap with other categories, but only 15 of 226 snowshoe prefixes had a second classification.

Each prefix was annotated with any ASNs listed in the SBL records as the "malicious ASN". To find these ASNs each SBL record was searched for any matches of the regular expression: 'AS[0-9]+' which matches on the string "AS" followed by one or more digits. ASNs were found for 190 (26.6%) of the DROP prefixes, the majority of which (130, 68.4%) Spamhaus classified as *hijacked*.

4.1.1 AFRINIC issues

Finally, 48.8% of the DROP address space related to two isolated AFRINIC incidents. In the first incident, an investigator claimed in 2019 that someone had manipulated the WHOIS records for several large prefixes managed by AFRINIC, allowing them to be sold fraudulently. An employee of AFRINIC allegedly exploited their position to modify the WHOIS database and then sold the affected IPv4 address space to a variety of Internet companies [69, 68]. 45 prefixes related to this incident appeared in the DROP list and were labeled *hijacked*.

The second incident was a legal dispute between AFRINIC and Cloud Innovations (AS398968). Cloud Innovation breached AFRINIC policy by deploying their AFRINIC resources outside of Africa. AFRINIC tried to reclaim these resources but Cloud Innovation has legally appealed this reclamation and so for now Cloud Innovation still holds the resources [71, 70]. Spamhaus has listed the 8 associated prefixes in DROP and labeled them *hijacked*, as they support the reclamation by AFRINIC.

Although the 45 prefixes related to both of these incidents made up only 6.3% of the prefixes in DROP, they made up 48.8% of the DROP address space. Their large size and anomalous character relative to the rest of the prefixes

labeled *hijacked* in DROP motivated us to exclude these 45 (of the 712 total DROP) prefixes from our analysis in further chapters to avoid skewing our characterization of the *hijacked* prefixes in DROP.

4.2 By RIR

In this section this work analyzed whether prefixes from a specific RIR were more heavily represented in DROP, and if so in which DROP category. This information could highlight systemic issues within an RIR leading to criminals targeting its address space for abuse.

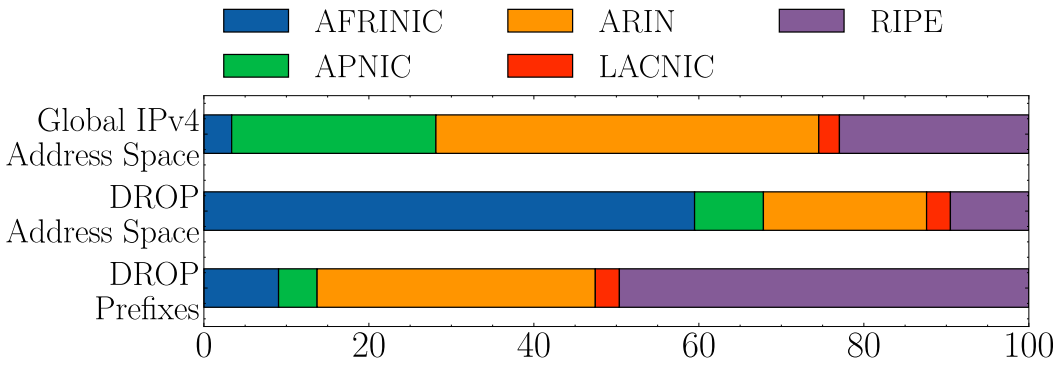

Figure 4.2: RIR managing prefixes in DROP (including prefixes related to AFRINIC issue discussed in §4.1.1). The top bar is the entire IPv4 address space by RIR (from the stats files §3.2.3). The bar below is the distribution of address space in DROP. RIPE has the largest percentage of prefixes but a much smaller percentage of address space in DROP. In contrast, AFRINIC has a small percentage of entries but a large percentage of the total address space covered by prefixes in DROP, because hijackers target large AFRINIC prefixes.

RIPE NCC allocated 46.8% of the prefixes in DROP, but these prefixes accounted for only 8.7% of the address space represented by all prefixes in DROP. In contrast, AFRINIC allocated 9% of prefixes in DROP but 57.7% of the DROP address space as illustrated in Figure 4.2. (AFRINIC prefixes comprise only 3.4% of the global IPv4 address space.) This difference is consistent with AFRINIC's smaller membership and delay in enforcement of IPv4 exhaustion policies (which reduced the maximum prefix allocation size) relative

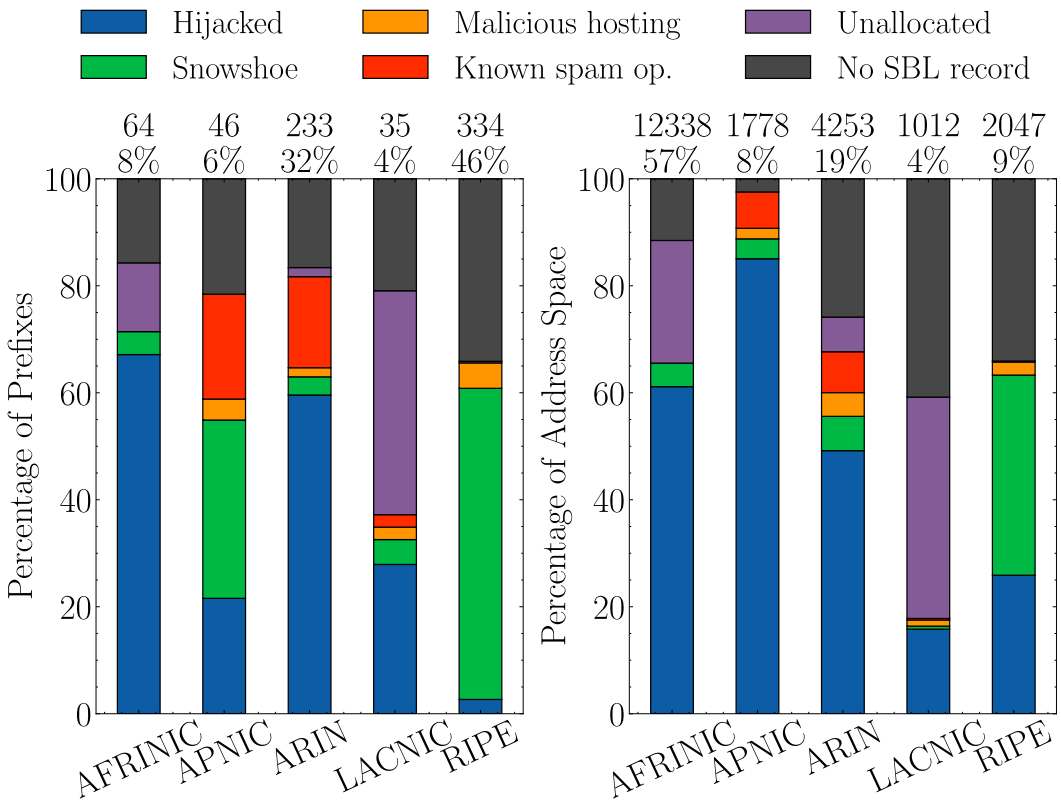

Figure 4.3: Breakdown of categories for Spamhaus DROP prefixes (left) and address space (right) by RIR. Prefixes from AFRINIC, APNIC, and ARIN were primarily hijacked prefixes, LACNIC had almost no hijacked prefixes but the majority of unallocated prefixes. RIPE NCC had the largest portion of suspected snowshoe prefixes.

to other RIRs as discussed in §2.3.

Figure 4.3 shows the distribution of DROP prefixes across RIRs as a function of DROP's abuse category. AFRINIC and ARIN had the highest percentage of hijacked prefixes: 47 of the 64 (73.4%) AFRINIC prefixes, and 141 of 231 (60.5%) ARIN prefixes. AFRINIC and ARIN have the lowest percentages of actively used address space – the largest fraction of address space that is allocated but not routed [74]. Prior work established that hijackers target these unused prefixes [72]. Further, in a recent study of the IPv4 address space, AFRINIC and ARIN had the lowest percentage of RPKI-signed prefixes [13]. In contrast, LACNIC and RIPE NCC had the highest percentage of actively used address space, and RPKI-signed prefixes [74, 13]. DROP labeled 5 (14.3%) of the 35 LACNIC prefixes, and 8 (2.4%) of the 333 RIPE NCC

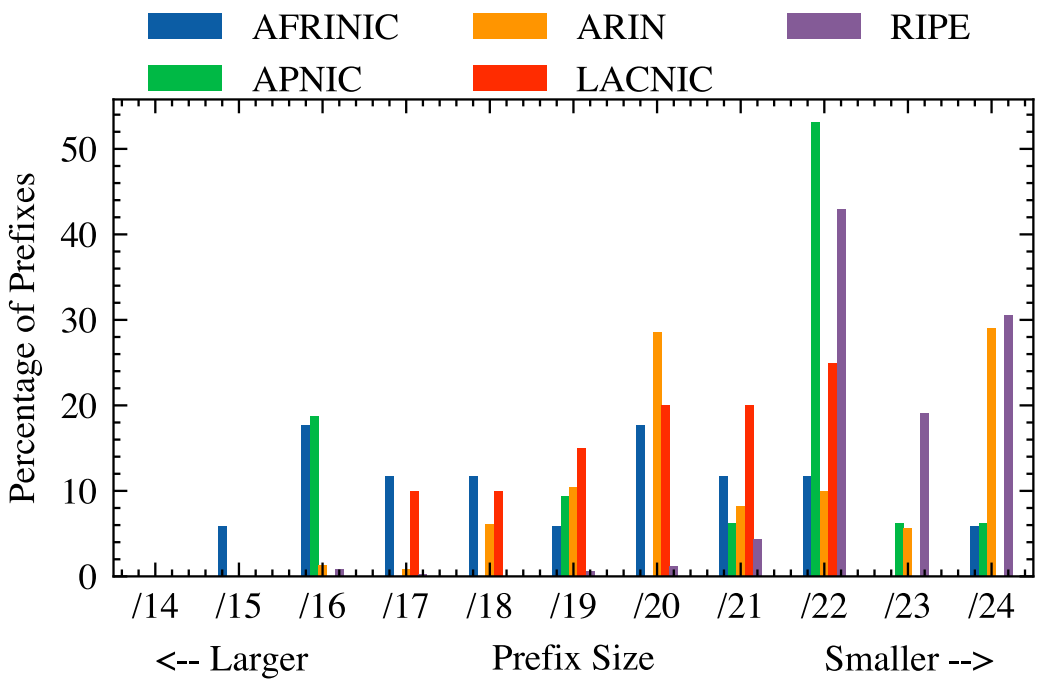

Figure 4.4: Size of prefixes in DROP by RIR. AFRINIC had large prefixes in DROP relative to the prefixes from other RIRs even after filtering out the anomalous AFRINIC prefixes (§4.1.1, which are not shown here) 60% of AFRINIC prefixes were larger than /20. In contrast, 98% of RIPE prefixes were /20 or smaller.

prefixes as *hijacked*.

RIPE NCC had the majority of *snowshoe spam* prefixes in DROP. There were 196 RIPE NCC prefixes classified as *snowshoe* by Spamhaus, representing 86.7% of the *snowshoe* prefixes. This may be related to criminals being less likely to face legal consequences in Europe [41].

4.2.1 Effect of Policies Inhibiting RPKI Uptake

As described in §2.4, ARIN and AFRINIC do not allow legacy resource holders to sign their legacy prefixes without signing a RSA. These two RIRs also had the largest number of legacy prefixes that appeared in DROP. 55 ARIN legacy prefixes appeared in DROP, and Spamhaus labeled 15 of these hijacked. Similarly, 41 AFRINIC legacy prefixes appeared in DROP, and Spamhaus labeled 35 of these hijacked. This was far more in comparison to the RIRs

that do allow legacy resource holders to sign legacy prefixes; 7 appeared from APNIC, 2 from RIPE and 1 from LACNIC. Thus, RIR policies may be having an observable impact on the ability of an attacker to conduct abuse on the address space managed by these RIRs, and increased RPKI deployment in other regions may create incentives for attackers to exploit address space that does not qualify for use of the RPKI.

4.3 By Allocation Status

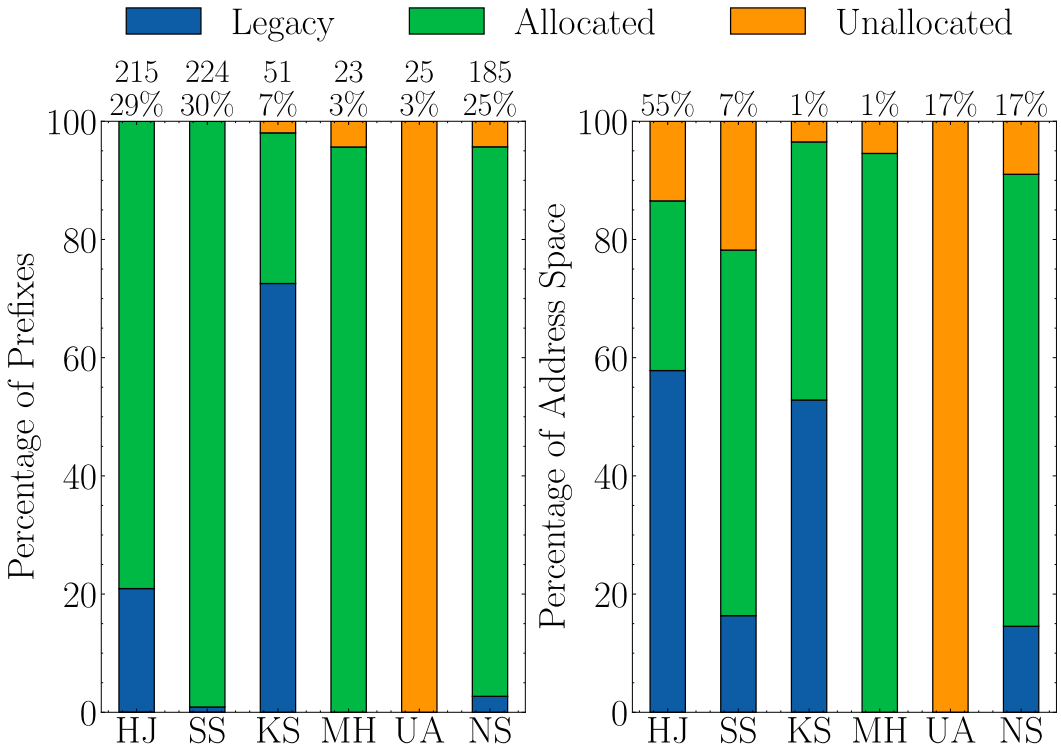

Figure 4.5: Allocation status of prefixes in DROP by abuse category. Hijackers are more likely to target legacy prefixes, 42% of hijacked address space was legacy as of March 2022, which is greater than the percentage (34%) of legacy prefixes in the total IPv4 address space. The many known spam operation legacy prefixes relate to a single operation that specializes in acquiring address space for spammers.

Figure 4.5 shows the allocation of prefixes and address space by category for the 712 prefixes in DROP. There were 30 legacy *known spam operation* prefixes (45% of *known spam operation* address space), all related to a single spam

operation that Spamhaus referred to as Big Sky Services, which specializes in acquiring address space and leasing it to spammers [62]. One method they used to acquire address space was purchasing the assets of defunct companies and then maintaining the appearance that the company was still in operation [62]. The abusers acquired an old legitimate ISP AS3502 (Intelligence Network Online, Inc.) that had legacy address space, and leased the prefixes to the spamming networks AS204472 and AS203999.

33 prefixes labeled hijacked (52% of hijacked address space) were legacy, a much larger percentage relative to the global legacy IPv4 address space, which was 34% in December 1997 [27], but is even less now due to legacy space being returned to RIRs. This shows that hijackers were more often targeting legacy resources for their hijacks, which was likely directly related to the large size of these prefixes, that they were often dormant, and for ARIN and AFRINIC because the resource holder had no means of deploying RPKI.

4.4 By BGP Activity

Samples of prefixes were taken from each DROP category and a large contrast in these prefixes BGP activity was observed proximate to their appearance in DROP. A timeline of the BGP activity for these samples is provided in Appendix A. The hijacked prefixes were mostly announced for short periods. Most hijacked prefixes either started being announced or had a change in origin AS within 2 months before they appeared in DROP which was likely the initiation of the hijack for malicious activity.

The snowshoe category contained more prefixes announced for longer periods, with no correlation to the date they appeared in DROP. However, there were still prefixes (e.g. 68.66.60.0/24) whose behavior resembles that of the hijacked prefixes. The ASes announcing known spam operation prefixes consistently announced them in BGP for long periods. Although some of these prefixes changed origin AS, there was no correlation between these changes

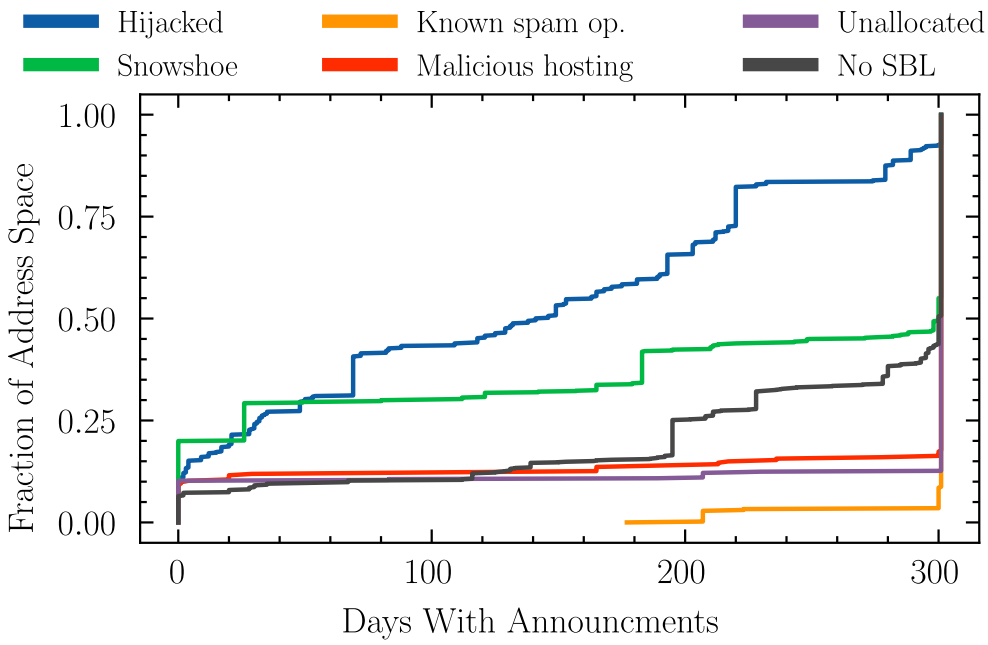

Figure 4.6: CDF for the number of days prefixes in DROP had at least one AS announcing them for the 300 days before they appeared in DROP. 81% of the hijacked address space was announced for at least one day but not the entire 300 days of our window, consistent with the hypothesis that hijackers are more often targeting dormant address space. Both the known spam operation and malicious hosting prefix categories had more than 80% of the corresponding address space announced on all of the 300 days, consistent with the hypothesis that this address space is legitimately allocated.

and when they appeared in DROP. The BGP activity for these prefixes was more comparable to the BGP activity of a legitimate prefix.

To examine whether the features we observed in these samples hold for all prefixes Figure 4.6, shows the duration of prefixes observed in BGP in the 300-day window before they appeared in DROP. 33.3% of hijacked address space was announced for between 1 and 100 days. Only 7% of hijacked address space was announced for all 200 days, confirming that most hijacked prefixes in DROP were previously dormant (unrouted by the legitimate owner). There were no previous announcements for 12.4% of the address space that DROP inferred as hijacked. Three scenarios could lead to this outcome: (1) Spamhaus added to DROP a less specific prefix of a hijacked prefix; (2) A prefix was stolen by means other than BGP hijacking; or (3) Spamhaus incorrectly inferred a

hijack.

Of the prefixes labeled as *snowshoe spam*, (30%) of the contained address space was not observed in BGP, and 55% was observed on all 300 days. Spamhaus will pre-emptively add address space to DROP if it has threat intelligence that suggests the address space will be used for the dissemination of spam [59], which explains the large percentage of address space categorized as *snowshoe* but not observed in BGP. The known spam operation and malicious hosting categories both had more than 80% of address space advertised in BGP on all of the 300 days before they appeared in DROP.

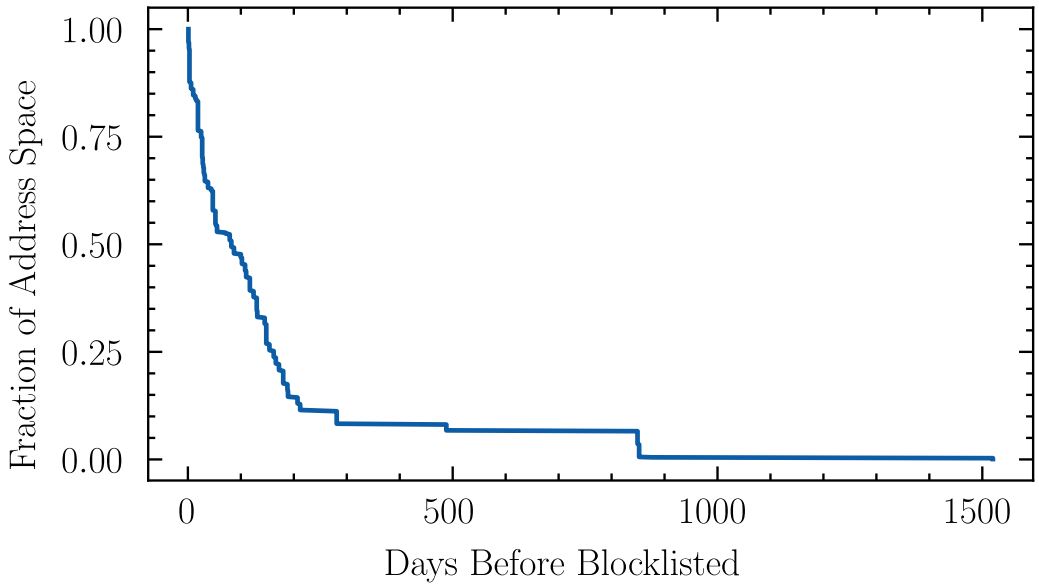

Figure 4.7: The number of days prior to appearing in DROP that each prefix labeled hijacked appeared in BGP originated by the SBL-labeled AS. 51% of hijacked address space did not appear in DROP until more than 2 months after the presumed hijacker announced it.

For each prefix in DROP labeled *hijacked* the hijack took place prior to the prefix appearing in DROP. This illustrates the limitation of blocklisting as a mitigation mechanism; it is reactive in nature. Figure 4.7 shows the time between each prefix first being announced by the hijacker in BGP and appearing in DROP. More than half of the hijacked address space did not appear in DROP until more than 2 months after it was hijacked. In this interim period the hijacker could have used the prefix maliciously and by the

time it is blocklisted the hijacker could already be targeting different address space. This creates a cat and mouse like situation where hijackers are always one step ahead of blocklist maintainers.

4.4.1 Forged LOA Hijacks

For a network operator to have a prefix routed by a transit provider, the operator may have to prove their ownership of the address space with a Letter of Authorization (LOA). Hijackers often forge these LOAs and illegitimately claim address space that does not belong to them.

On July 1, 2021, four /24 prefixes owned by four different defunct companies appeared in DROP. The SBL record for each prefix stated that a hijacker took control of the prefix by forging an LOA. In this case, the hijacker registered the domain name for each of the defunct businesses, providing false credibility that they were the legitimate owner of the address space. No ASes were observed advertising these prefixes from January 2015 until March 2022, so it is unlikely that the hijacker ever used this address space.

Address space was also discovered that exhibited BGP activity resembling a forged LOA hijack. Two large transit provider ASes were observed originating BGP advertisements for address space that Spamhaus labeled hijacked. Large transit providers are not likely to intentionally hijack address space, what is more likely is that they were deceived into advertising this address space on behalf of a malicious actor. Of the ASes that were observed originating announcements for DROP prefixes in the 200-day window before the prefix appeared in DROP, Figure 4.8 shows that only AS174 (Cogent) and AS3356 (Level3) had a customer cone size (the size of the set of ASes that can be reached from a given AS following only customer links) larger than 10,000 ASes. In total ,these two ASes originated announcements for 19 large DROP prefixes which composed 28.7% of the hijacked DROP address space.

Between July 2019 and September 2019, AS174 illegitimately announced 6 /16 prefixes that later appeared in DROP. Two of these prefixes were Aus-

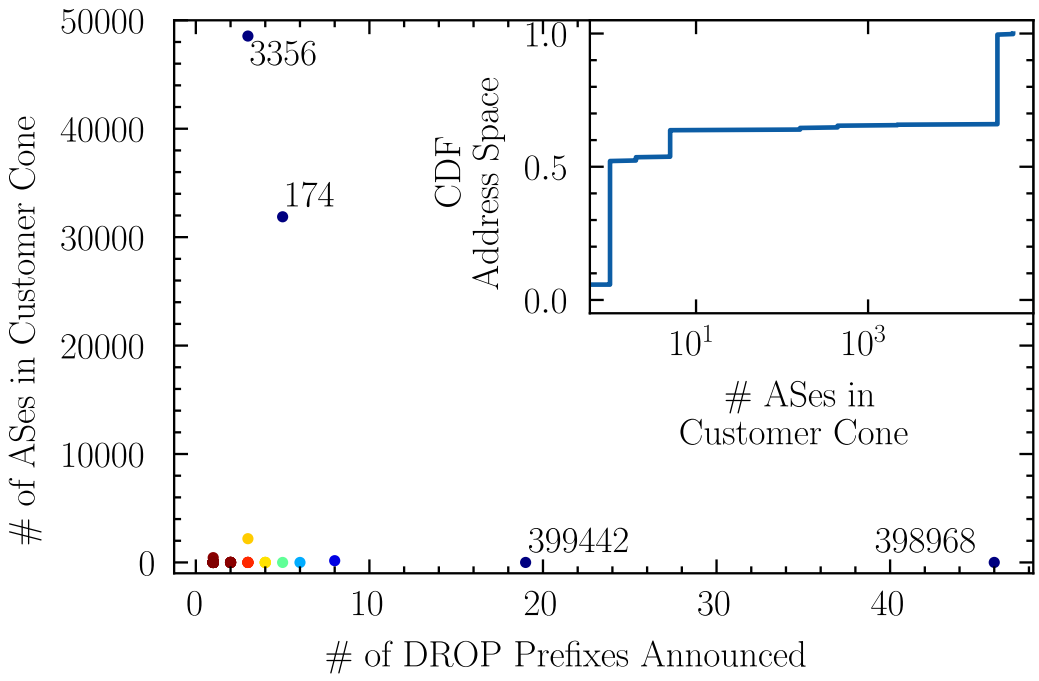

Figure 4.8: Each origin AS in BGP announcements for prefixes in DROP. There is a large concentration of small ASes observed announcing 1-20 DROP prefixes. The larger ASes observed were likely deceived by malicious actors into announcing the prefixes.

tralian government networks. From June 2021 to July 2021, Level3 (AS3356) advertised five prefixes that Spamhaus labeled hijacked. Both AS174 and AS3356 withdrew the prefixes from BGP shortly after the prefixes appeared in DROP, presumably after learning the illegitimate nature of the announcements.

4.4.2 Collateral Prefixes

In some instances when an AS hijacked a prefix, Spamhaus added all other prefixes announced by that AS to DROP despite some of them being legitimately allocated to the hijacking AS.

On January 28, 2020, 8 prefixes appeared in DROP that AS18013 had announced. Each SBL record stated ‘AS18013 ASLINE - IP hijackers’, however, AS18013 controlled the WHOIS record for 3 of the prefixes. Spamhaus informed us that when they identified AS18013 as a hijacker, Spamhaus added

all prefixes announced by AS18013 to DROP.

On February 2, 2020, Spamhaus added to DROP all of the address space (4 prefixes) announced by the actor allegedly behind the AFRINIC hijacks (Netstyle A. LTD, AS199267, AS43945, AS58018). One of these prefixes was a hijacked AFRINIC prefix but the other three were legitimately allocated to Netstyle.

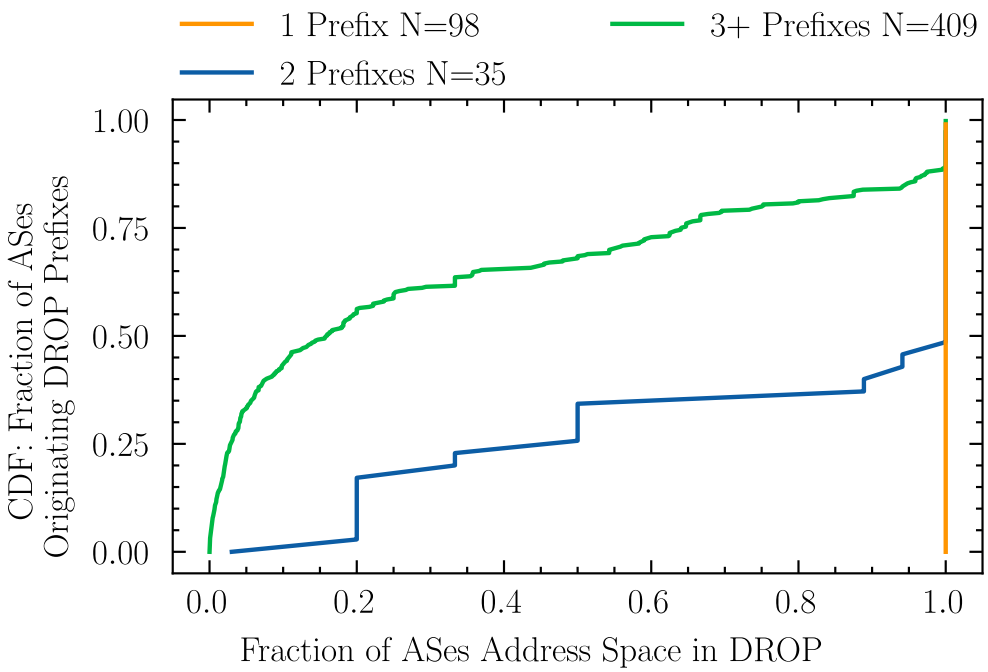

Figure 4.9: Fraction of an AS's address space in DROP on the date its prefix(es) appeared in DROP. For ASes that only announced one prefix, Spamhaus always added the entire prefix. 11% of ASes originating 3 or more prefixes in BGP had all of their address space in DROP. In these instances it is possible that Spamhaus blocklisted the entire AS's routed address space irrespective of it being used maliciously or not.

Figure 4.9 shows the fraction of an AS's address space that was in DROP at the time each AS's prefix appeared in DROP. The 11% of ASes with 3 or more prefixes in BGP that had their entire address space in DROP at the time their prefix appeared in DROP were examined. In addition to ASLINE and Netstyle, 2 more ASes that had all of their address space in DROP with SBL records that said they were a hijacker when the prefix was legitimately allocated to the AS were found. In total of the prefixes that appeared in

DROP, 11 were found to be collateral prefixes, which was 2.2% of the hijacked address space.

Chapter 5

Effect of Blocklisting

This section, investigates the effect that blocklisting had on routing visibility and RPKI uptake.

5.1 Routing Visibility

The available BGP data set suggests that a prefix being listed in DROP caused some of the attackers (or their transit providers) to withdraw the prefix as shown in Figure 5.1: 19% of the prefixes were withdrawn within 30 days of being listed in DROP. For prefixes that Spamhaus labeled *hijacked* or *unallocated* the percentage was higher: 70.7% and 54.8%, respectively. These two categories stand out as expected; prefixes in these categories were being advertised illegitimately, and illegitimate announcers likely withdrew prefixes once the addresses were less effective for their malicious applications. Further, at the transit level, 4 RouteViews peers provided full tables but appeared to use the DROP list to filter prefixes listed in DROP from BGP announcements (right panel in Figure 5.1). The categories with low fractions of prefixes that were withdrawn from BGP contained mostly prefixes that RIRs legitimately allocated to these ASes who were using them maliciously, e.g. bulletproof hosting companies.

The category with the largest percentage of address space that was deallocated by an RIR after appearing in DROP was *malicious hosting*. Of the *malicious*

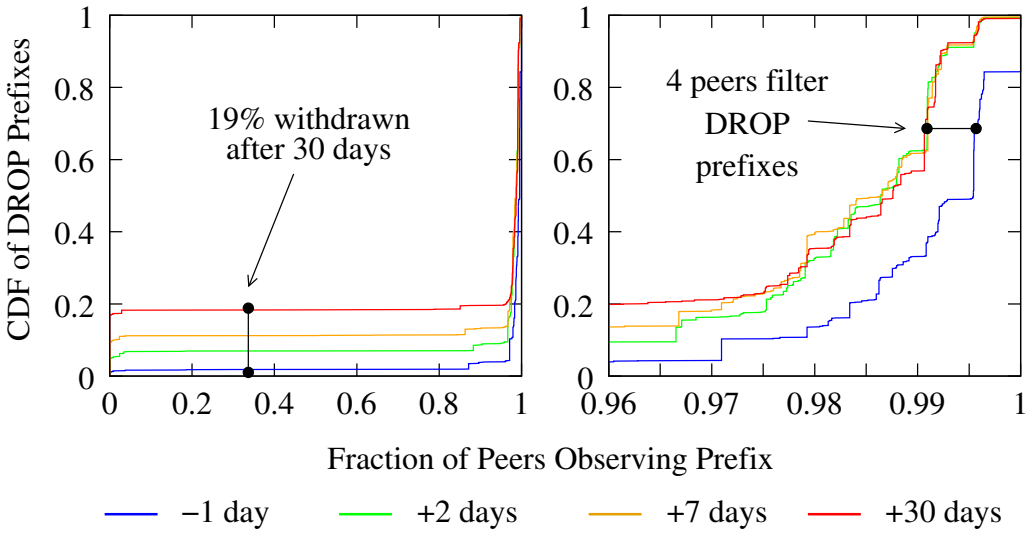

Figure 5.1: This plot shows visibility of prefixes (fraction of peers observing each DROP prefix) on the x axis, and a CDF of the fraction of total DROP prefixes on the y axis. When comparing the visibility before the prefixes appear in DROP to after they appear in DROP, there is a reduction in visibility. 19% of prefixes listed in DROP were not BGP-observed 30 days after being listed, suggesting that appearing in DROP may have motivated the malicious actor to withdraw the prefix from BGP. Four RouteViews peers appeared to filter DROP prefixes from their BGP announcements, suggesting that some networks use DROP to filter in the (BGP) control plane.

hosting prefixes, 17.4% were already allocated when they appeared in DROP and had their allocations withdrawn by the end of March 2022. A similar pattern occurred for prefixes Spamhaus removed from DROP: 8.8% of the prefixes Spamhaus removed were deallocated. Half of these prefixes Spamhaus removed within a week of an RIR deallocating them, which suggests that Spamhaus removed them as a result of their deallocation.

5.2 Improved RPKI Uptake

Table 5.1 examines the RPKI properties of 650 prefixes that were not RPKI-signed when they were added to DROP. Each region has a base level of RPKI activity, with the RIPE region having the largest fraction of unsigned prefixes whose resource holders signed during this study. For all but AFRINIC, the

	Never in DROP	Removed from DROP	Present in DROP
AFRINIC	11.8% of 3901	14.3% of 7	0.0% of 11
APNIC	26.3% of 42.2K	44.4% of 18	21.6% of 37
ARIN	8.5% of 65.2K	25.0% of 40	0.6% of 169
LACNIC	25.5% of 15.1K	35.1% of 37	0% of 9
RIPE NCC	33.0% of 68.2K	54.2% of 83	19.8% of 172
Overall	22.3% of 195.6K	42.5% of 186	13.8% of 420

Table 5.1: RPKI signing rate of prefixes without a ROA that Spamhaus added to DROP between June 5, 2019 and March 30, 2022. A larger fraction of prefixes had a ROA created if they were removed from DROP than if they were never in DROP for most regions. A smaller fraction of prefixes had a ROA created if they were not removed from DROP.

signing rate of prefixes that Spamhaus removed from DROP was larger than this base level, indicating that operators were motivated to deploy RPKI as part of the process of getting their prefix removed from DROP. This process requires the prefix owner to prove that the problem has been resolved and will not continue [61]. Security-conscious network operators most likely own these prefixes, and will more likely deploy preventive measures to prevent their re-appearance in DROP. Similarly, the signing rate of prefixes that remained in DROP was lower than this base level, suggesting that these prefixes are relatively neglected.

Chapter 6

Effectiveness of IRR

This chapter, investigates the effectiveness of IRR by analyzing how operators and attackers both used IRR for the prefixes that appeared in DROP.

In the 7-day window before appearing in DROP, 226 DROP prefixes (31.7%) covering 68.8% of the DROP address space had either a route object in the RADb IRR (§2.2.1) with an exact match or a more specific prefix. The RADb IRR contains evidence suggesting attackers used IRR to make their activities appear legitimate, as 32% of the prefixes with route objects had their route object created during the month before the prefix appeared in DROP. More encouragingly, 43% of prefixes with route objects had their route object removed a month after they appeared in DROP.

Focusing on the 130 DROP prefixes whose SBL record reported the prefix as hijacked by a specific ASN, 69 (55%) either had no route object or had a route object with a different AS from the DROP-labeled AS. The remaining 57 prefixes (45%) had the labeled hijacking ASN in the route object, suggesting the hijacker was able to circumvent the RADb authorization process. Overall, there were 13 different hijacking ASNs in the route objects for these 57 prefixes.

Of the 57 prefixes, 49 had route objects with different origin ASes but shared three different ORG-IDs, indicating the bulk of the fraudulent entries were created by three entities. One of these ORG-IDs created IRR records for 15 of the hijacked DROP prefixes, using various origin ASes. Each of these

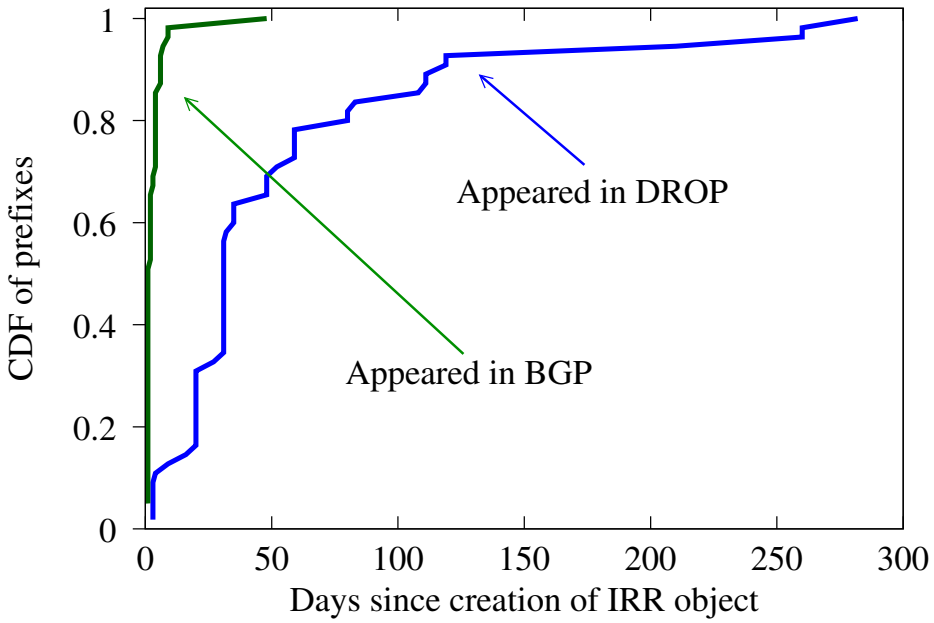

Figure 6.1: Other than 2 prefixes that had already been in BGP for over a year, the remaining prefixes appeared in BGP shortly after creation of the IRR record.

prefixes were announced shortly after the route object was created and each had a common AS in their announced path: AS50509, which hijacked unrouted prefixes using defunct ASes as the origin, and creating IRR route objects with these defunct ASes. For all but 2 of the 57 prefixes, the hijacker AS started announcing the prefix in BGP less than a week after they created the IRR record (Figure 6.1). In the other 2 instances, the hijacker created the IRR record more than a year after they had already been announcing the prefix in BGP. The prefixes these hijacking ASes were targeting were all apparently abandoned, as only 5 of the prefixes had existing IRR entries prior to the DROP-labeled AS creating their entry.

There was 1 prefix in DROP that was unallocated when an AS created a route object for it. No unallocated prefix should be routed, therefore unallocated prefixes should not be accepted into the IRR. The fact that it was further highlights the lack of verification performed by RADb.

These findings prove that IRR databases in their current form can not be relied on to accurately authenticate BGP announcements. To remediate this

problem IRR operators need to implement more rigorous ownership validation processes in which they ensure that route objects may only be created by entities who are legitimately allocated each specific IPv4 resource. This could be achieved in collaboration with RIRs, as RIRs hold the information regarding the ASes who have been allocated IPv4 resources. Another takeaway from these findings and possibility for future work is that the creation of forged IRR records could be used as a preemptive signal to identify BGP hijacks, seeing as in most instances attackers created the IRR records just days before they hijacked the prefix.

Chapter 7

Effectiveness of RPKI

This chapter, (1) investigates the effectiveness of RPKI by analyzing hijacks of any RPKI-signed prefixes added to DROP, (2) discusses the implications of these hijacks for RPKI-signed prefixes more broadly, and (3) discusses potential solutions such as AS0.

7.1 Evidence of RPKI-valid Hijacks

Of the 179 prefixes labeled hijacked, only three were RPKI-signed *before* they were blocklisted. This infers that hijackers do not usually target RPKI-signed prefixes but rather target unallocated or unrouted non-RPKI signed address spaces. The entity allegedly hijacking two of these prefixes appeared to control the ROA, as the ASN in published ROAs changed when the BGP origin ASN changed in the two years prior to the prefix appearing in DROP. However, the third prefix is a real-world demonstration of the limitations in capability of the current RPKI deployment.

Consider the hijacked RPKI-signed prefix 132.255.0.0/22 illustrated in Figure 7.1, with a ROA authorizing AS263692 — a Peruvian network under LACNIC — to originate the prefix. While AS263692 routed the prefix via a South American transit provider (AS21575) for many years, in July 2020 it stopped i.e. , the prefix became unrouted. In December 2020, we see the prefix again originated by AS263692 but routed via Russian ASes AS50509 and AS34665

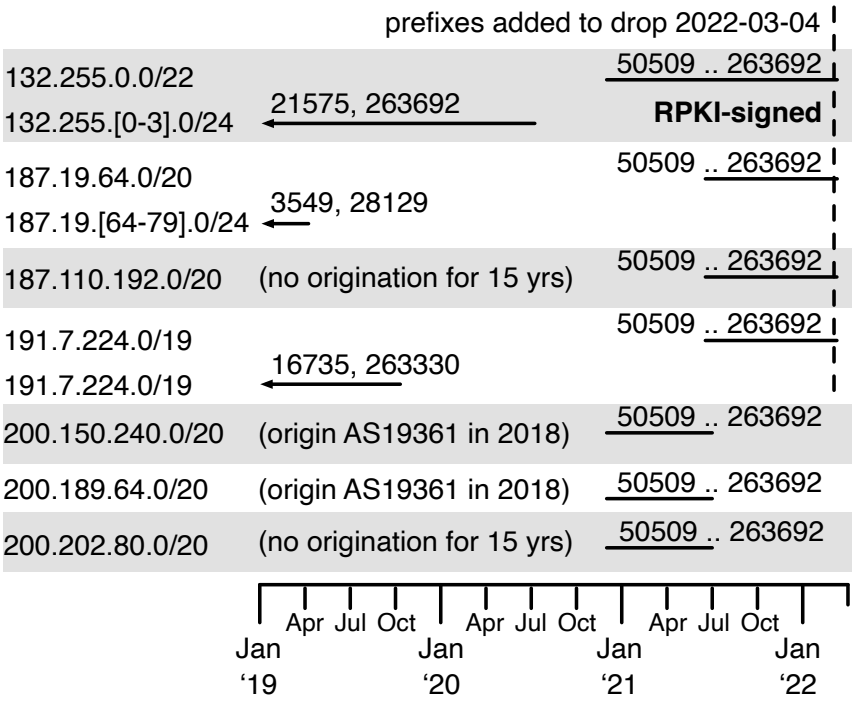

Figure 7.1: Case study of hijacker considering origin AS of historic announcements. AS263692 is a Peruvian AS with historic transit through a South American transit provider (AS 21575) and a single RPKI-signed prefix: 132.255.0.0/22. In December 2020, a hijacker begins announcing that prefix, along with prefixes historically unrouted or originated with a different ASN, with further announcements in June 2021, through a Russian transit provider (AS 50509).

who had hijacked the prefix. Recall, AS50509 is also implicated in hijacking unrouted prefixes by creating IRR route objects with the defunct ASes (Chapter 6). Since the origin AS matched the ROA, the announcement was deemed by validating software as RPKI-valid, subverting RPKI protections. On inspecting the BGP routing data for a similar pattern — originated by AS263692 routed via AS50509 — six additional non-RPKI signed prefixes were found (Figure 7.1). Of these six, only three prefixes were added to DROP by Spamhaus.

To prevent hijacks of unrouted RPKI-signed prefixes, the ROA should use AS0 as the authorized origin. As such, the underlying reason why the hijack of 132.255.0.0/22 was successful was because not only did the AS not route the prefix, but the ROA contained a non-AS0 ASN. Effectively, this hijack implies

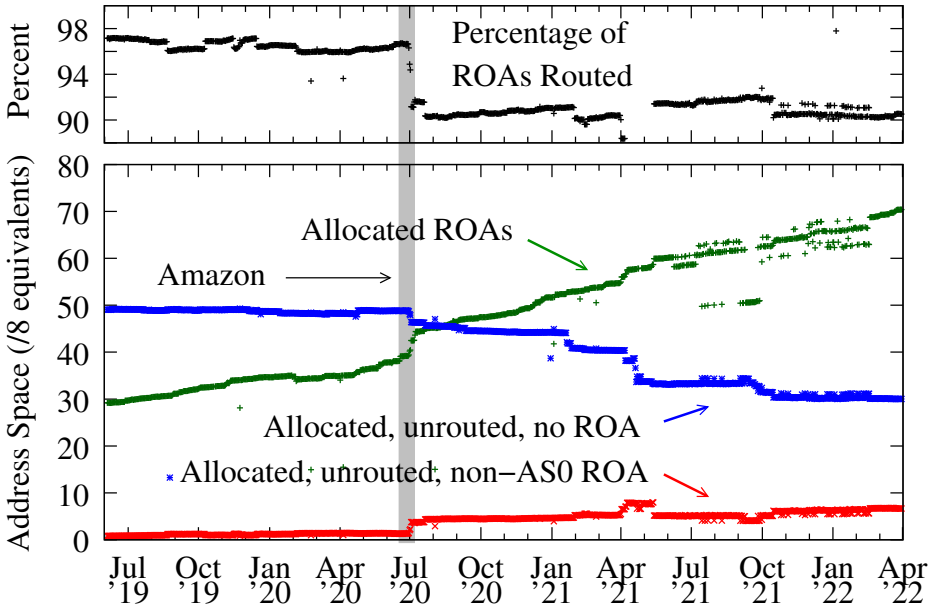

Figure 7.2: Routing status of ROAs. The majority of ROAs were routed, though the percentage of signed address space routed declined from 97.1% to 90.5%. Address space equivalent to 6.7 /8s was signed with a non-AS0 ROA and unrouted as of March 2022.

that any *unrouted* non-AS0 RPKI-signed prefixes are no better protected than non-RPKI signed prefixes. Figure 7.2 shows that while the amount of IPv4 address space covered by a ROA has increased, the volume of unrouted but signed prefixes has also increased, and as of March 2022, the equivalent of 6.7 /8 prefixes (≈112M IPs) are signed but not routed. While these prefixes are susceptible to hijacks, this risk could be eliminated by the holder signing the ROAs with AS0.

Figure 7.2 also shows that as of March 2022, the equivalent of 30.0 /8s were allocated but unrouted and had no ROA. The RIRs managing this address space were examined, and it was found that the equivalent of 18.25 /8s (60.8%) was managed by ARIN. Because ARIN manages the bulk of this allocated but unrouted address space, we encourage ARIN members to develop policy that incentivizes resource holders to not only use RPKI but also issue AS0 ROAs when they have no plans to announce the prefix.

7.2 AS0 Policies at Operator and RIR level

An AS0 ROA prevents unallocated or unrouted address space from being routed [24] (Section 2.2.2). Given the potential for AS0 policies to considerably reduce the attack surface in today's routing system, this section discusses the different policy considerations and challenges at the operator and RIR level.

7.2.1 Operator AS0

While there was the equivalent of 6.7 /8s unrouted and covered by a non-AS0 ROA (Figure 7.2) the bulk of this address space (70.1%, the equivalent of 4.7 /8s) was held by three organizations: Amazon with the equivalent of 3.1 /8s (ROA creation event labeled in Figure 7.2), Prudential Insurance with one unrouted /8, and Alibaba with the equivalent of 0.64 /8s. As such, a few organizations adopting AS0 could remediate the majority of the attack surface.

Notably, one DROP prefix was RPKI-signed with an AS0 ROA by a network operator. Spamhaus added 45.65.112.0/22 to DROP on January 28, 2020. The operator signed it with AS0 on May 5, 2021. Spamhaus removed it from DROP on June 16, 2021, suggesting that the AS0 ROA played a part in its removal.

The most likely reason that the unannounced hijacked address space that appeared in DROP never got signed with AS0 is because the address space was abandoned with no one to sign it. Another reason operators may hesitate to RPKI-sign their unused address space with AS0 is because it indicates to RIRs that address space is not being used, and RIRs are actively seeking to reclaim IP address space.

7.2.2 RIR AS0

RIRs can create AS0 ROAs for unallocated prefixes. In 2019, APNIC was the first RIR to propose an *AS0 policy*. The policy was controversial: some critics

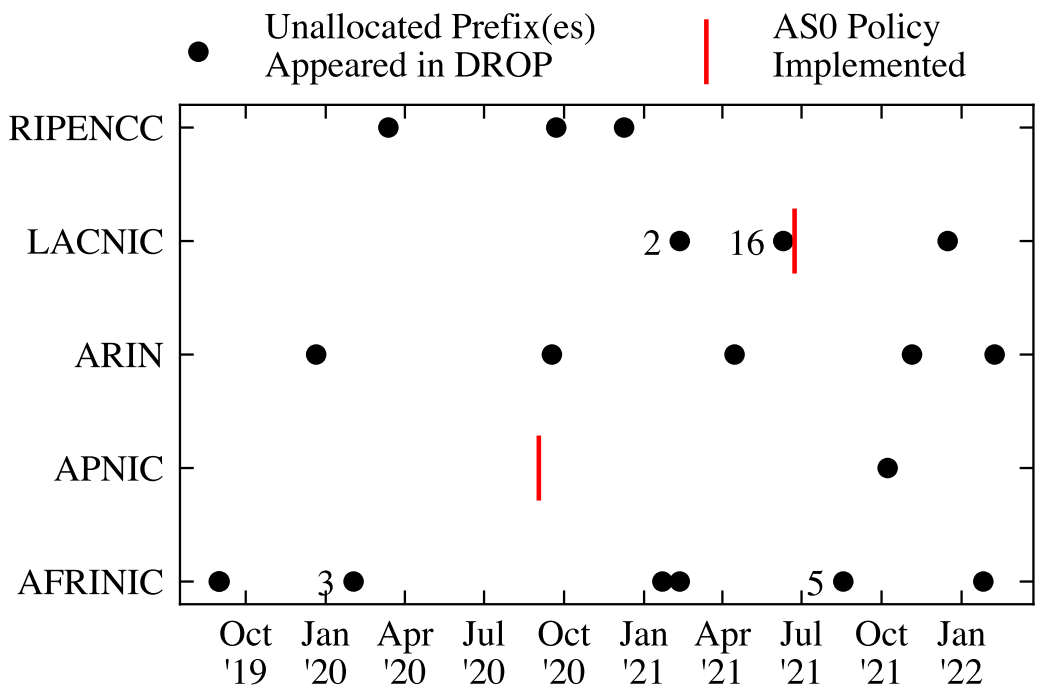

Figure 7.3: Timeline of when address space unallocated by RIRs appeared in DROP, and when an AS0 policy was implemented by an RIR. In practice, the ability of an attacker to hijack an unallocated prefix is not affected by the RIRs' current AS0 policies.

considered it a dangerous slippery slope giving the RIRs too much operational responsibility when they were not 24x7 operations [19]. APNIC implemented an AS0 policy [5] on September 2, 2020, and LACNIC implemented an AS0 policy on June 23, 2021 [48]. RIPE NCC proposed an AS0 policy on October 22, 2019, but ultimately withdrew the proposal [2]. AFRINIC proposed an AS0 policy in November 2019, but have yet to implement it as of May 2022 [22]. ARIN has not made any AS0 proposal.

From the period of June 5, 2019 to March 30, 2022, 40 unallocated prefixes appeared in DROP, with events clustered for LACNIC (19) and AFRINIC (12) resources. The size of these clusters is not correlated with the amount of unallocated address space remaining in the RIRs (Figure 7.4). Hijacks of unallocated address space continued beyond the implementation of an AS0 policy, due to (1) the RIRs implementing their AS0 policy using a different Trust Anchor Locator (TAL) (file used to allow relying parties to retrieve RPKI

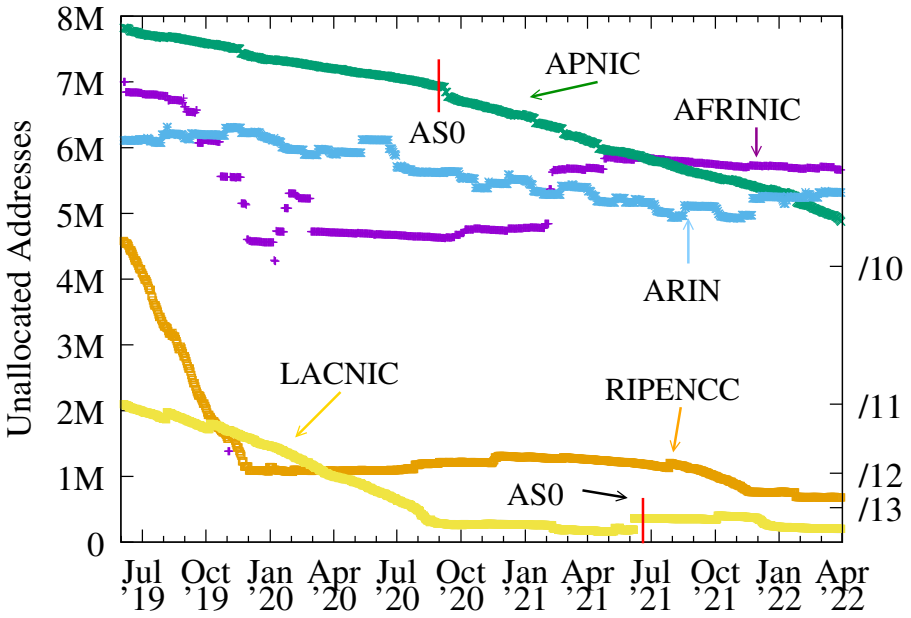

Figure 7.4: Amount of unallocated address space remaining in each RIR's free pool, over time.

data from a repository) that is not configured in any RPKI validation software by default, and (2) recommend solely for information purposes, rather than for route filtering. Of the RouteViews tables for peers that provided a full routing table on March 30, 2022, none appeared to use APNIC or LACNIC AS0 TALs to filter routes, as every peer reported ≈30 prefixes that would have been filtered with those TALs.

Chapter 8

Conclusion

This work used 712 prefixes from the last three years of Spamhaus' DROP list as a lens to analyze IP address abuse risks and mitigations. It was found that the vast majority of hijacked prefixes that were listed in DROP had been left dormant by their owners. An RPKI AS0 ROA would have mitigated these hijacks, however, often the address space was allocated to a defunct organization so there was no one with authority to deploy RPKI. This introduces the question should RIRs have the power to intervene in these situations and either deploy an AS0 ROA or reclaim this address space? Another finding of this work was that blocklisting had an effect on the 712 prefixes that were added to DROP – for a hijacked or unallocated prefix being added to the DROP list led most attackers to withdraw those routes, and prefixes that were blocklisted were more likely to adopt RPKI than prefixes that were not. This research also found evidence of a vulnerability in RPKI that illustrates the hijack risk to *all* unrouted RPKI-signed prefixes, equivalent to 6.7 /8s, ≈112M IPs. While unrouted RPKI-signed prefixes can use AS0 ROAs to prevent hijacks, the unrouted unsigned prefixes (equivalent to 30.0 /8s, ≈480M IPs) will continue to be easy targets for hijackers. As such, these results indicate that policies concerning RPKI, and AS0 more specifically, merit re-evaluation by both operators and RIRs.

Appendices

A

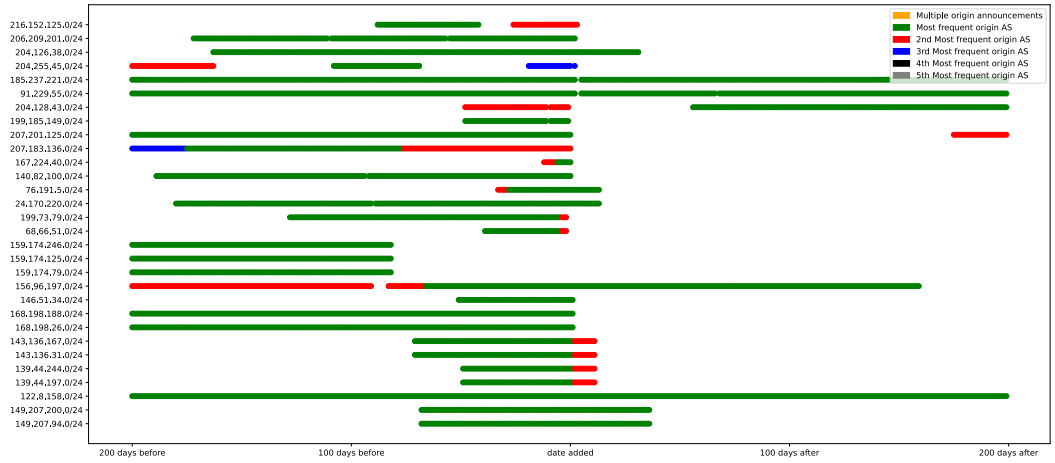

Figure .1: BGP announcements for a random sample of 30 hijacked prefixes within a 400 day window. Prefixes are announced for short periods and often announcements start shortly before appearing in the blocklist. We stop observing announcements for a significant portion of prefixes once they appear in the blocklist.

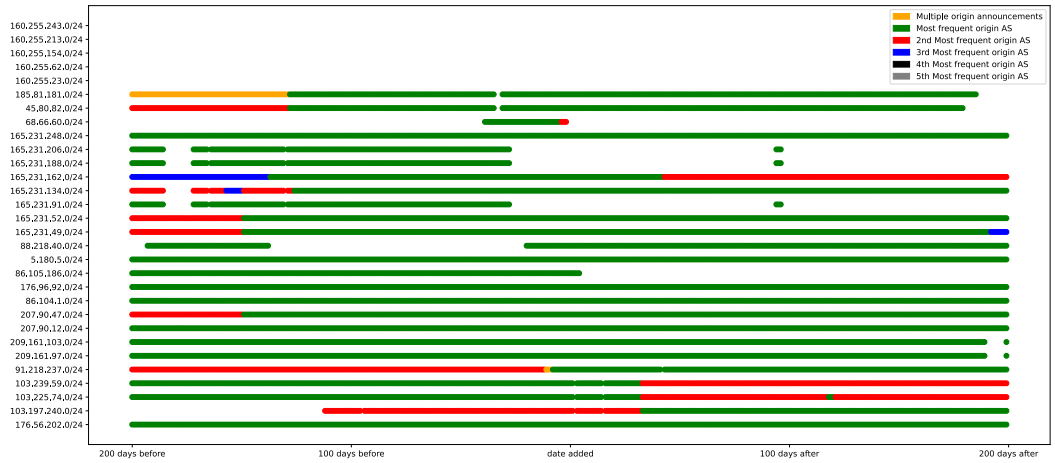

Figure .2: BGP announcements for a random sample of 30 snowshoe prefixes within a 400 day window. Overall prefixes are announced for longer and more consistently than the hijacked prefixes. There are fewer prefixes with correlation between announcements and the date the prefix was blocklisted, in comparison to hijacked prefixes. There are still some prefixes e.g. 68.66.60.0/24 which resemble a hijacked prefix.

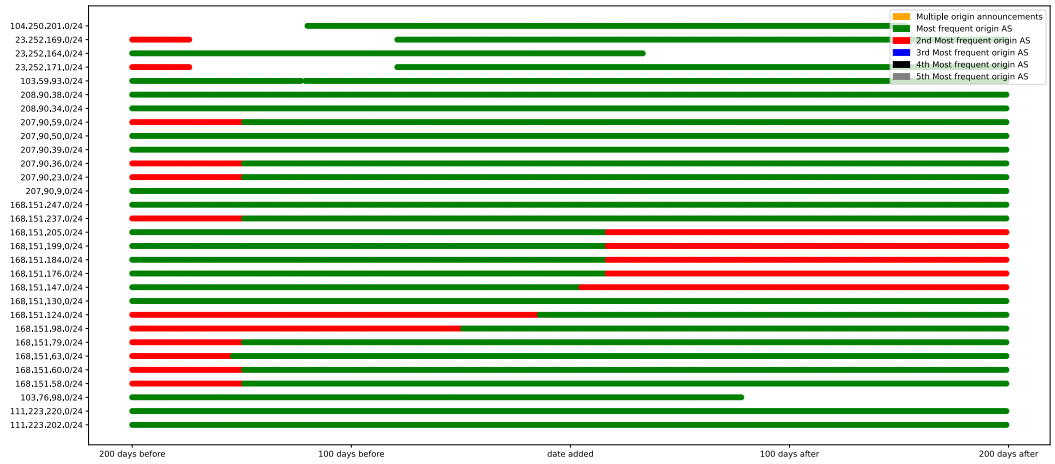

Figure .3: BGP announcements for a random sample of 30 known spam operation prefixes within a 400 day window. Prefixes are announced for long periods consistently, and there is no correlation with the date the prefix was blocklisted. There is no evidence of hijacks, instead they resemble legitimately routed prefixes. The prefixes which are synchronised (change origin at the same time) are prefixes which are being announced by the same entities, as a large portion of the known spam operations are associated with a single actor.

References

[1] Greg Aaron, Lyman Chapin, David Piscitello, and Dr. Colin Strutt. Whois contact data availability and registrant classification study: A study of the effects of gdpr and icann policy, 2020.

[2] Melchior Aelmans, Martijn Schmidt, and Massimiliano Stucchi. Slurm file for unallocated and unassigned ripe ncc address space, Oct 2019.

[3] AFRINIC. IPv4 exhaustion phase 2. `https://afrinic.net/20200113-afrinic-enters-ipv4-exhaustion-phase-2`, January 2020.

[4] AFRINIC. Legacy resource holders, Nov 2020.

[5] APNIC. Prop-132: RPKI ROAs for unallocated and unassigned APNIC address space (was: AS0 for Bogons). `https://www.apnic.net/community/policy/proposals/prop-132`, 2020.

[6] APNIC. RIR statistics exchange format. `https://www.apnic.net/about-apnic/corporate-documents/documents/resource-guidelines/rir-statistics-exchange-format/`, 2022.

[7] Maria Apostolaki, Aviv Zohar, and Laurent Vanbever. Hijacking Bitcoin: Routing Attacks on Cryptocurrencies. In *2017 IEEE Symposium on Security and Privacy (SP)*, pages 375–392, San Jose, CA, USA, May 2017. IEEE.

[8] ARIN. Services available to organizations holding legacy resources, 2022.

[9] Rob Austein, Steven Bellovin, Russ Housley, Stephen Kent, Warren Kumari, Doug Montgomery, Chris Morrow, Sandy Murphy, Keyur Patel, John Scudder, Samuel Weiler, Matthew Lepinski, and Kotikalapudi Sriram. RFC 8205 - BGPsec protocol specification. In *IETF RFCs*, page 45. 2017.

[10] Hitesh Ballani, Paul Francis, and Xinyang Zhang. A study of prefix hijacking and interception in the internet. *SIGCOMM Comput. Commun. Rev.*, 37(4):265–276, aug 2007.

[11] Xavier Le Bris. Status of Legacy IPv4 Address Space. https://labs.ripe.net/author/xavier/status-of-legacy-ipv4-address-space/, 2011.

[12] Shinyoung Cho, Romain Fontugne, Kenjiro Cho, Alberto Dainotti, and Phillipa Gill. BGP hijacking classification. In *TMA*, pages 25–32, 2019.

[13] 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 *IMC*, pages 406–419, Amsterdam Netherlands, October 2019. ACM.

[14] Catalin Cimpanu. Google traffic hijacked via tiny Nigerian ISP. https://www.zdnet.com/article/google-traffic-hijacked-via-tiny-nigerian-isp/, November 2018.

[15] Catalin Cimpanu. Russian telco hijacks internet traffic for Google, AWS, Cloudflare, and others. https://www.zdnet.com/article/russian-telco-hijacks-internet-traffic-for-google-aws-cloudflare-and-others/, April 2020.

[16] California Civil Code. California consumer privacy act, 2018. Resolution AB-375 (2017–2018 Session) Website Assembly Bill No. 375.

[17] Leslie Daigle. WHOIS protocol specification. RFC 3912, September 2004.

[18] Alun Davies. Outcome of the afrinic audit, 2022.

[19] Owen DeLong. prop-132-v002: As0 for bogons, Aug 2019.

[20] Ben Du, Gautam Akiwate, Thomas Krenc, Cecilia Testart, Alexander Marder, Bradley Huffaker, Alex C Snoeren, and KC Claffy. Irr hygiene in the rpki era. In *International Conference on Passive and Active Network Measurement*, pages 321–337. Springer, 2022.

[21] Dan Goodin. Suspicious event hijacks Amazon traffic for 2 hours, steals cryptocurrency, April 2018.

[22] Frank Habicht, Mark Elkins, Jordi Palet Martinez, and Haitham El Nakhal Hytham. Rpki roas for unallocated and unassigned afrinic address space (draft 3), Mar 2022.

[23] Xin Hu and Z. Morley Mao. Accurate Real-time Identification of IP Prefix Hijacking. In *IEEE S&P*, pages 3–17, May 2007.

[24] Geoff Huston and George G. Michaelson. Validation of route origination using the resource certificate public key infrastructure (PKI) and route origin authorizations (ROAs). RFC 6483, February 2012.

[25] Geoff Huston, Mattia Rossi, and Grenville Armitage. Securing BGP — A literature survey. *IEEE Communications Surveys Tutorials*, 13(2):199–222, 2011.

[26] IANA. Internet assigned numbers authority, 2022.

[27] IANA. IPv4 address space registry. `https://www.iana.org/assignments/ipv4-address-space/ipv4-address-space.xhtml`, 2022.

[28] IRR. Internet routing registry. `https://www.irr.net/`, 2022.

[29] Varun Khare, Qing Ju, and Beichuan Zhang. Concurrent prefix hijacks: Occurrence and impacts. In *IMC*, pages 29–36, November 2012.

[30] Christopher Kruegel, Darren Mutz, William Robertson, and Fredrik Valeur. Topology-Based Detection of Anomalous BGP Messages. In Giovanni Vigna, Christopher Kruegel, and Erland Jonsson, editors, *Recent Advances in Intrusion Detection*, Lecture Notes in Computer Science, pages 17–35, Berlin, Heidelberg, 2003. Springer.

[31] Warren Kumari, Barry Leiba, Suzanne Woolf, Joe Abley, Tim April, Paul Ebersman, Ondrej Filip, Geoff Huston, Jacques Latour, John Levine, Chris Roosenraad, and Tara Whalen. Sac109 - the implications of dns over https and dns over tls, 2020.

[32] LACNIC. Lacnic legacy resources.

[33] LACNIC. Lacnic rpki.

[34] LACNIC, 2022.

[35] Mohit Lad, Dan Massey, Dan Pei, Yiguo Wu, Beichuan Zhang, and Lixia Zhang. PHAS: A prefix hijack alert system. In *USENIX Security*, July 2006.

[36] Matt Lepinski and Stephen Kent. An infrastructure to support secure internet routing. RFC 6480, February 2012.

[37] Vector Guo Li, Gautam Akiwate, Kirill Levchenko, Geoffrey M. Voelker, and Stefan Savage. Clairvoyance: Inferring Blocklist Use on the Internet. In Oliver Hohlfeld, Andra Lutu, and Dave Levin, editors, *PAM*, volume 12671, pages 57–75. Springer International Publishing, Cham, 2021.

[38] Vector Guo Li, Matthew Dunn, Paul Pearce, Damon McCoy, Geoffrey M Voelker, Stefan Savage, and Kirill Levchenko. Reading the Tea Leaves: A Comparative Analysis of Threat Intelligence. page 18.

[39] Marc Lindsey. Protect your pre-1997 IP address. `https://www.computerworld.com/article/2514777/protect-your-pre-1997-ip-address.html`, December 2010.

[40] Ratul Mahajan, David Wetherall, and Tom Anderson. Understanding BGP misconfiguration. *ACM SIGCOMM Computer Communication Review*, 32(4):3–16, August 2002.

[41] Stephen McCombie, Josef Pieprzyk, and Paul Watters. *Cybercrime Attribution: An Eastern European Case Study*. 2009.

[42] Merit Network. The Internet Routing Registry - RADb, 2021.

[43] Asya Mitseva, Andriy Panchenko, and Thomas Engel. The state of affairs in bgp security: A survey of attacks and defenses. *Computer Communications*, 124:45–60, 2018.

[44] NIST. RPKI monitor. `https://rpki-monitor.antd.nist.gov/`, 2022.

[45] Ostap Efremov. 196.52.0.0/14 revoked, cleanup efforts needed. RIPE NCC Anti-Abuse Working Group, 2021.

[46] Pierluigi Paganini. BGP hijacking - Traffic for Google, Apple, Facebook, Microsoft and other tech giants routed through Russia. `https://securityaffairs.co/wordpress/66838/hacking/bgp-hijacking-russia.html`, December 2017.

[47] European Parliament and Council of the European Union. General data protection regulation, May 2016. Journal Reference L119.

[48] Ricardo Patara and Aftab Sidiqui. Rpki asn 0 roa policy, May 2020.

[49] Jian Qiu, Lixin Gao, Supranamaya Ranjan, and Antonio Nucci. Detecting bogus BGP route information: Going beyond prefix hijacking. In *2007 Third International Conference on Security and Privacy in Communications Networks and the Workshops - SecureComm 2007*, pages 381–390, September 2007.

[50] RADb. RADb Archive. ftp://ftp.radb.net, 2022.

[51] Anirudh Ramachandran and Nick Feamster. Understanding the network-level behavior of spammers. 36(4):291–302, aug 2006.

[52] RIPE. Resource certification (rpki) for provider independent end users and legacy end users, Feb 2021.

[53] RIPE. RIPE RPKI Archive. https://ftp.ripe.net/ripe/rpki/, 2022.

[54] RIPE. RIPE stats FTP repository. https://ftp.ripe.net/pub/stats/, 2022.

[55] Ronald F. Guilmette. Cogent & FDCServers: Knowingly aiding and abetting fraud and theft?, 2019.

[56] Sanjaya. Rpki services now available to apnic historical resource holders, Mar 2021.

[57] Pavlos Sermpezis, Vasileios Kotronis, Petros Gigis, Xenofontas Dimitropoulos, Danilo Cicalese, Alistair King, and Alberto Dainotti. ARTEMIS: Neutralizing BGP Hijacking Within a Minute. *IEEE/ACM Transactions on Networking*, 26(6):2471–2486, December 2018.

[58] Aftab Siddiqui. KlaySwap - another BGP hijack targeting crypto wallets. https://www.manrs.org/2022/02/klayswap-another-bgp-hijack-targeting-crypto-wallets/, February 2022.

[59] Spamhaus.

[60] Spamhaus. DROP - Don't Route or Peer. https://www.spamhaus.org/drop/, 2022.

[61] Spamhaus. DROP (FAQ). https://www.spamhaus.org/faq/section/Spamhaus\%C2\%A0DROP, 2022.

[62] Spamhaus. ROKSO: Big Sky Services / Corespace / Mark Wulff / Liana Dunlap. https://www.spamhaus.org/rokso/evidence/ROK11571/big-sky-services-corespace-mark-wulff-liana-dunlap/main-info, 2022.

[63] Spamhaus. SBL - Spamhaus blocklist. https://www.spamhaus.org/sbl/, 2022.

[64] Meenakshi Syamkumar, Ramakrishnan Durairajan, and Paul Barford. Bigfoot: A geo-based visualization methodology for detecting BGP threats. In *2016 IEEE Symposium on Visualization for Cyber Security (VizSec)*, pages 1–8, Baltimore, MD, USA, October 2016. IEEE.

[65] Cecilia Testart, Philipp Richter, Alistair King, Alberto Dainotti, and David Clark. Profiling BGP Serial Hijackers: Capturing Persistent Misbehavior in the Global Routing Table. In *IMC*, pages 420–434, Amsterdam Netherlands, October 2019. ACM.

[66] Andree Toonk. Looking at the spamhaus ddos from a bgp perspective, Mar 2013.

[67] Costa Tsaousis. FireHOL IP Lists — IP Blacklists — IP Reputation Feeds. `http://iplists.firehol.org/`, 2022.

[68] Jan Vermeulen. The big South African IP address heist – How millions are made on the "grey" market. `https://mybroadband.co.za/news/internet/318205-the-big-south-african-ip-address-heist-how-millions-are-made-on-the-grey-market.html`, September 2019.

[69] Jan Vermeulen. How Internet resources worth R800 million were stolen and sold on the black market. `https://mybroadband.co.za/news/internet/330379-how-internet-resources-worth-r800-million-were-stolen-and-sold-on-the-black-market.html`, December 2019.

[70] Jan Vermeulen. Afrinic bank accounts frozen after R740 million damages claim. `https://mybroadband.co.za/news/internet/407770-afrinic-bank-accounts-frozen-after-r740-million-damages-claim.html`, July 2021.

[71] Jan Vermeulen. Internet addresses worth R1.8 billion seized. `https://mybroadband.co.za/news/internet/405640-internet-addresses-worth-r1-8-billion-seized.html`, July 2021.

[72] Pierre-Antoine Vervier, Olivier Thonnard, and Marc Dacier. Mind Your Blocks: On the Stealthiness of Malicious BGP Hijacks. In *NDSS*, San Diego, CA, 2015. Internet Society.

[73] Matthias Wählisch, Olaf Maennel, and Thomas Schmidt. Towards detecting bgp route hijacking using the rpki. *ACM SIGCOMM Computer Communication Review*, 42:103–104, 10 2012.

[74] Sebastian Zander, Lachlan L. H. Andrew, Grenville Armitage, and Geoff Huston. Estimating IPv4 address space usage with capture-recapture. In *38th Annual IEEE Conference on Local Computer Networks - Workshops*, pages 1010–1017, Sydney, Australia, October 2013. IEEE.